User's Guide

Notice

Enterasys Networks reserves the right to make changes in specifications and other information contained in this document without prior notice. The reader should in all cases consult Enterasys Networks to determine whether any such changes have been made.

The hardware, firmware, or software described in this manual is subject to change without notice.

IN NO EVENT SHALL ENERASYS NETWORKS BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING BUT NOT LIMITED TO LOST PROFITS) ARISING OUT OF OR RELATED TO THIS MANUAL OR THE INFORMATION CONTAINED IN IT, EVEN IF ENERASYS NETWORKS HAS BEEN ADVISED OF, KNOWN, OR SHOULD HAVE KNOWN, THE POSSIBILITY OF SUCH DAMAGES.

IN NO EVENT WILL ENTERASYS NETWORKS BE LIABLE FOR ANY DAMAGES WHATSOEVER, INCLUDING SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THE WEB LOG VIEW, WHETHER IN AN ACTION OF CONTRACT OR TORT INCLUDING NEGLIGENCE.

Virus Disclaimer

Enterasys Networks has tested its software with current virus checking technologies. However, because no anti-virus system is 100% reliable, we strongly caution you to write protect and then verify that the Licensed Software, prior to installing it, is virus-free with an anti-virus system in which you have confidence.

Enterasys Networks makes no representations or warranties to the effect that the Licensed Software is virus-free.

Copyright © 2000, 2001, 2002 by Enterasys Networks. All rights reserved.

Printed in the United States of America.

Order Number: 9032722-05 August 2002

Enterasys Networks, Inc.
P.O. Box 5005
Rochester, NH 03866-5005

Enterasys, NetSight, XPedition and **Matrix E7** are trademarks of Enterasys Networks.

Windows, Windows NT, Windows XP, Microsoft, Microsoft Windows, and **Microsoft Windows for Workgroups** are trademarks or registered trademarks of Microsoft Corporation.

Solaris, Sun, Sun Microsystems, are trademarks or registered trademarks of Sun Microsystems, Inc.

Cabletron Systems, **SPECTRUM**, **BRIM**, **DNI**, **FNB**, **INA**, **Integrated Network Architecture**, **LANVIEW**, **LANVIEW Secure**, **Multi Media Access Center**, **MiniMMAC**, and **TRMM** are registered trademarks, and **Bridge/Router Interface Modules**, **BRIM-A100**, **CRBRIM-W/E**, **CRXMIM**, **CXRMIM**, **Desktop Network Interface**, **Distributed LAN Monitoring**, **Distributed Network Server**, **DLM**, **DNSMIM**, **E1000**, **E2000**, **E3000**, **EFDMIM**, **EMM-E6**, **EMME**, **EPIM**, **EPIM-3PS**, **EPIM-A**, **EPIM-C**, **EPIM-F1**, **EPIM-F2**, **EPIM-F3**, **EPIM-T**, **EPIM-T1**, **EPIM-X**, **ESXMIM**, **ETSMIM**, **ETWMIM**, **FDCMIM-04**, **FDCMIM-08**, **FDMMIM**, **FDMMIM-04**, **Flexible Network Bus**, **FOMIM**, **FORMIM**, **HubSTACK**, **IRBM**, **IRM**, **IRM-2**,

IRM-3, **Media Interface Module**, **MicroMMAC**, **MIM**, **MMAC**, **MMAC-3**, **MMAC-3FNB**, **MMAC-5**, **MMAC-5FNB**, **MMAC-8**, **MMAC-8FNB**, **MMAC-M8FNB**, **MMAC-Plus**, **MRX**, **MRXI**, **MRXI-24**, **MultiChannel**, **NB20E**, **NB25E**, **NB30**, **NB35**, **NBR-220/420/620**, **RMIM**, **SecureFast Switch**, **SecureFast Packet Switching**, **SFS**, **SFPS**, **SPECTRUM Element Manager**, **SPECTRUM for Open Systems**, **SPIM-A**, **SPIM-C**, **SPIM-F1**, **SPIM-F2**, **SPIM-T**, **SPIM-T1**, **TPMIM**, **TPMIM-22**, **TPMIM-T1**, **TPRMIM**, **TPRMIM-36**, **TPT-T**, **TRBMIM**, **TRMM-2**, **TRMMIM**, and **TRXI** are trademarks of Cabletron Systems, Inc.

AppleTalk, Apple, Macintosh, and TokenTalk are registered trademarks; and Apple Remote Access and EtherTalk are trademarks of Apple Computer, Inc.

SmartBoost is a trademark of American Power Conversion

ST is a registered trademark and C++ is a trademark of AT&T

Banyan and VINES are registered trademarks of Banyan Systems, Inc.

cisco, ciscoSystems, and AGS+ are registered trademarks; and cBus, cisco Router, CRM, IGS, and MGS are trademarks of cisco Systems, Inc.

GatorBox is a registered trademark; and GatorMail, GatorMIM, GatorPrint, GatorShare, GatorStar, GatorStar GX-M, and XGator are trademarks of Cayman Systems, Inc.

CompuServe is a registered trademark of CompuServe Incorporated

X Window System is a trademark of Consortium, Inc.

CTERM, DECnet, and ULTRIX are registered trademarks; and DEC, DEC C++, DECnet-DOS, DECstation, VAX DOCUMENT, VMA, and VT are trademarks of Digital Equipment Corporation

Fore Systems, ForeRunner, and ForeRunner ASX-100 are trademarks of Fore Systems, Inc.

PC/TCP is a registered trademark of FTP Software, Inc.

HP OpenView is a registered trademark of Hewlett-Packard, Inc.

AIX, IBM, OS/2, NetView, and PS/2 are registered trademarks; and AT, Micro Channel, PC, PC-DOS, PC/XT, Personal Computer AT, Operating System/2, Personal System/2, RISC System/6000, and Workplace Shell are trademarks of International Business Machines Corporation

i960 microprocessor is a registered trademark; and Intel and Multichannel are trademarks of Intel Corporation

Microsoft, MS-DOS, and Windows are registered trademarks of Microsoft Corporation

Chameleon, ChameleonNFS, Chameleon 32, IPX/link, and NEWT are trademarks of NETMANAGE, Inc.

NetWare and Novell are registered trademarks; and Internetwork Packet Exchange (IPX), IPX, and Network File System (NFS) are trademarks of Novell, Inc.

Motif and MS are registered trademarks; and Open Software Foundation, OSF, OSF/1, and OSF/Motif are trademarks of The Open Software Foundation, Inc.

Silicon Graphics and IRIS are registered trademarks; and Indigo and IRIX are trademarks of Silicon Graphics, Inc.

NFS, PC-NFS, SPARC, Sun Microsystems, and Sun Workstation are registered trademarks; and OpenWindows, SPARCstation, SPARCstation IPC, SPARCstation IPX, Sun, Sun-2, Sun-3, Sun-4, Sun386i, SunNet, SunOS, SunSPARC, and SunView are trademarks of Sun Microsystems, Inc.

OPEN LOOK and UNIX are registered trademarks of Unix System Laboratories, Inc.

Ethernet, NS, Xerox Network Systems and XNS are trademarks of Xerox Corporation

ANNEX, ANNEX-II, ANNEX-IIe, ANNEX-3, ANNEX-802.5, MICRO-ANNEX-XL, and MICRO-ANNEX-ELS are trademarks of Xylogics, Inc.

MAXserver and Xyplex are trademarks of Xyplex, Inc.

Restricted Rights Notice

(Applicable to licenses to the United States Government only.)

1. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

   Enterasys Networks, 35 Industrial Way, Rochester, New Hampshire 03867-0505.

2. (a) This computer software is submitted with restricted rights. It may not be used, reproduced, or disclosed by the Government except as provided in paragraph (b) of this Notice or as otherwise expressly stated in the contract.

   (b) This computer software may be:

   (1) Used or copied for use in or with the computer or computers for which it was acquired, including use at any Government installation to which such computer or computers may be transferred;

   (2) Used or copied for use in a backup computer if any computer for which it was acquired is inoperative;

   (3) Reproduced for safekeeping (archives) or backup purposes;

   (4) Modified, adapted, or combined with other computer software, provided that the modified, combined, or adapted portions of the derivative software incorporating restricted computer software are made subject to the same restricted rights;

   (5) Disclosed to and reproduced for use by support service contractors in accordance with subparagraphs (b) (1) through (4) of this clause, provided the Government makes such disclosure or reproduction subject to these restricted rights; and

   (6) Used or copied for use in or transferred to a replacement computer.

   (c) Notwithstanding the foregoing, if this computer software is published copyrighted computer software, it is licensed to the Government, without disclosure prohibitions, with the minimum rights set forth in paragraph (b) of this clause.

(d) Any other rights or limitations regarding the use, duplication, or disclosure of this computer software are to be expressly stated in, or incorporated in, the contract.

(e) This Notice shall be marked on any reproduction of this computer software, in whole or in part.

Contents

Chapter 1 Introduction

Using This Guide 1-1
Getting Help 1-2
- Using On-line Help 1-2
- Accessing On-line Documentation 1-3
  - Documentation Feedback 1-3
- Getting Technical Support 1-3
  - Online Services on the World Wide Web 1-3
  - Global Technical Assistance Center 1-3

Chapter 2 Overview of NetSight Element Manager

About NetSight Element Manager 2-1
- Network Discovery 2-4
- Management and Monitoring Workspace 2-5
  - List Views 2-6
  - Tree Views 2-7
  - Map Views 2-9
  - Log Views 2-10
- Excluding Nodes 2-10
- Alarm and Event Handling 2-11

The NetSight Element Manager Tool Suite 2-11
- Remote Administration Tools 2-12
- Scheduler Tool 2-12
- MIB Tools 2-13

Customizing NetSight Element Manager 2-13
- Adjusting Default Settings 2-13
  - Workspace Options 2-14
  - Network Options 2-15
  - Polling Options 2-17
  - Node Options 2-20
  - Directories Options 2-21
  - Device Management Options 2-23
  - Event Log Options 2-24
- Customizing the Toolbar 2-26
  - Adding a Tool Entry 2-27
  - Editing a Tool Entry 2-30
  - Removing a Tool Entry 2-30

Backing Up Data Files 2-30

Editing the Backup Configuration File 2-33
Performing a Backup 2-34
Restoring Files from a Backup Set 2-35
Restoring Default Files 2-35

Chapter 3 Discovering Nodes

Accessing the Discover Manager Window 3-1
Creating and Modifying Discover Scripts 3-2
Accessing the Discover Properties Window 3-2
Setting General Discover Properties 3-4
Setting Device Discovery Parameters 3-5
Defining IP Address Ranges 3-6
Defining and Selecting Subnets 3-7
Defining a Community Name 3-11
Selecting a Script Filter 3-12
Setting a Discover Interval 3-13
Executing Discover Scripts 3-14
The Discover Process 3-16
NEW and SECONDARY Discovers 3-17
Stopping a Discover in Progress 3-17
Restarting a Completed Discover 3-17
Scheduled Discovers 3-18
Deleting a Discover Script 3-18
Starting the Discover Wizard 3-19
Launching the Scheduler Application 3-19
Using the IP or Host Locator Tool 3-19
Matching a Host Name to an IP Address 3-20
Matching an IP Address to a Host Name 3-21
Locating the MAC Address for a Host Name/IP Pair 3-21
Subnet Search 3-22
Matching a MAC Address to a Host Name and IP 3-25
Using the Subnets Window 3-27
Creating and Modifying a Subnet Entry 3-28
Deleting a Subnet Entry 3-29

Chapter 4 List Views

About List Views 4-1
The List View Workspace 4-1
Adding a Folder to the Workspace 4-2
The All Nodes List View 4-3
Other Default List Views 4-4
List View Display Options 4-5
Using the List View Toolbar 4-5
Sorting Nodes 4-6
Searching a List View 4-7
Creating and Editing List Views 4-10

Selecting Node Information to Display 4-11
Node Information Options 4-12
Selecting a Display Font 4-14
Selecting the Nodes to Display: Using Filters 4-15
Filter Categories 4-16
Creating the Filter 4-18
Excluded Nodes 4-18
Excluding Nodes from an Individual View 4-18
To Exclude a Node from a View 4-19
To Re-include a Node in a View 4-21
Excluding Nodes from the Central Node Database 4-22
The Excluded Nodes List and Discover 4-23
To Exclude a Node from the Central Node Database 4-23
To Re-include a Node into the Central Node Database 4-24
To Delete a Node from the Database 4-25
Adding Network Elements to the Node Database 4-26
Adding a Device 4-27
Editing Device Properties 4-29
Configuring Node Interface Poll Settings 4-34
Editing Poll Settings for Multiple Nodes 4-36
Editing Node Class Properties 4-37
Selecting a New Model, Class, Enterprise, or Topology Symbol 4-39
Changing Management Application Information 4-40
Changing General Node Class Properties 4-44
List Views and Maps 4-46
Selecting and Deleting Orphans 4-46
Using the Go To Map Function 4-47

Chapter 5 Tree Views

About Tree Views 5-1
The Tree View Workspace 5-2
Adding a Folder to the Workspace 5-2
The All Nodes Tree View 5-4
Other Default Tree Views 5-5
Navigating a Tree View 5-5
Using the Toolbar and the View Menu 5-5
Using the Keyboard 5-6
Creating and Editing Tree Views 5-7
Selecting Node Information to Display 5-8
Node Properties Options 5-9
Interface Properties Options 5-11
Selecting a Display Font 5-12
Selecting the Nodes to Display: Using Filters 5-13
Filter Categories 5-13
Creating the Filter 5-15
Excluded Nodes 5-16
Tree Views and Maps 5-16
Using the Go To Map Function 5-16

Chapter 6 Creating Network Maps

Working with Maps ... 6-1
About the Map View ... 6-2
About Submap Icons ... 6-3
About GoTo Symbols ... 6-4
Map and Device Status Conditions ... 6-4
Building Your Network Map ... 6-6
Adding a Folder ... 6-6
Renaming a Folder ... 6-7
Adding a New Map ... 6-7
Setting Basic Map Options ... 6-9
Editing Map Properties ... 6-12
Adding a Submap ... 6-13
Generating Submaps Automatically ... 6-15
Changing a Submap Link ... 6-17
Adding a GoTo Symbol ... 6-18
Importing an MRM Map ... 6-20
Setting MRM Import Options ... 6-20
Importing a Map ... 6-21
Adding Network Elements to your Map ... 6-23
Adding a Device Directly to a Map ... 6-23
Customizing Your Map ... 6-26
Using Non-manageable Symbols ... 6-27
Placing a Symbol Icon ... 6-27
Editing Symbol Properties ... 6-28
Using Connections ... 6-33
Using Lines and Shapes ... 6-34
Using Text ... 6-37
Adding Object Notes ... 6-39
Adding Notes to Multiple Nodes ... 6-42
Changing the Layout of Map Objects ... 6-42
Layering Items in a Map ... 6-43
Align Left, Right, Top, or Bottom ... 6-43
Space Evenly (Horizontally or Vertically) ... 6-44
Make Same Width, Height, or Size ... 6-45
Align Objects in Grid ... 6-46
Align Objects in Ring ... 6-46
Using Object Linking and Embedding (OLE) ... 6-47
Embedding an Object in a Map File ... 6-49
Creating an Object Link ... 6-52
Adding an Object Package ... 6-52
Editing Object Data ... 6-53
Updating a Link ... 6-54
Converting an Object ... 6-56
Working with Linked and Embedded Objects ... 6-57
Editing Your Map ... 6-58
Protecting Your Maps from Editing: The Lock Feature ... 6-60
Managing Maps and Folders ... 6-61
Opening and Closing Maps and Folders ... 6-61

Saving Maps 6-62
Deleting Maps and Folders 6-62

Chapter 7 Node Management Overview

Device Management 7-1
Launching Device Management 7-1
A Brief Overview of Device Management 7-2
Source Address Naming 7-4
PING 7-6
Accessing the PING Window 7-6

Appendix A Batch Customization via AddTool, AddImage, and AddDevType

Using the AddTool Utility ..... Appendix A-1
Creating an Add Tool Input File ..... Appendix A-3
Comma-Separated Variable (CSV) Format ..... Appendix A-3
Tag/Data Format ..... Appendix A-3
Executing the AddTool Command ..... Appendix A-4
Using the AddImage Utility ..... Appendix A-5
Creating an AddImage Input File ..... Appendix A-6
Comma-Separated Variable (CSV) Format ..... Appendix A-6
Tag/Data Format ..... Appendix A-7
Executing the AddImage Command ..... Appendix A-8
Using the AddDevType Utility ..... Appendix A-9
Creating an AddDevType Input File ..... Appendix A-11
Comma-Separated Variable (CSV) Format ..... Appendix A-11
Tag/Data Format ..... Appendix A-12
Executing the AddDevType Command ..... Appendix A-12
Viewing the Log File ..... Appendix A-13

Chapter 1

Introduction

Using this guide; getting help

Using This Guide

This manual contains general information regarding NetSight Element Manager, including:

- Chapter 1, Introduction, which introduces the product, provides an outline of this book's structure, describes how to access and use the on-line documentation, and tells you how to get assistance while using the product.

- Chapter 2, Overview of NetSight Element Manager, provides a more in-depth look at some of the application's features and enhancements, and describes some general functions of the software platform such as configuring options, backing up data files, customizing the toolbar, and printing.

- Chapter 3, Discovering Nodes, provides details on how to use the Discover application to discover nodes on your network. It describes how to build scripts that will discover IP devices by address range or subnetwork range. The IP or Host Locator tool — which resolves IP addresses to host names, and vice versa — is also described here.

- Chapter 4, List Views, covers the List View workspace and List View windows. NetSight Element Manager stores all node information in a central database. List Views provide you with a variety of user configurable views into that database, including a wide selection of status and descriptive information about each node. This chapter also provides information on manually adding nodes to the node database, editing a node's displayed properties and individual polling intervals, and launching device management.

- Chapter 5, Tree Views, describes the Tree View workspace and windows. These user-configurable views let you display your network nodes in a hierarchical tree control format that provides a variety of information about the features supported by the device, including information about the interface through which NetSight Element Manager is communicating with each node.

- Chapter 6, Creating Network Maps, provides detailed instructions on how to use the Map View workspace, and how to create and link individual map files. You can use maps to represent your network configuration via symbols and icons. This chapter also describes the use of Microsoft Object Linking and Embedding (OLE) with map files.

- Chapter 7, Node Management Overview, explains several features of device management, including an introduction to the Device View windows; creating a source address definition file; and a brief introduction to the Microsoft PING application.

- Appendix A, **Batch Customization via AddTool, AddImage, and AddDevType**, Using the AddTool, AddImage, and AddDevType tools to expand the capabilities of NetSight Element Manager

Other manuals in your documentation set — including the device-specific User's Guides, Tools Guide, Remote Administration Tools User's Guide, Alarm and Event Handling Guide, and RMON User's Guide — contain specific information about the various applications included in your NetSight Element Manager platform. A set of installation instructions, a Getting Started guide, and a README.TXT file are also provided.

Getting Help

This section describes two different methods of getting help for questions or concerns you may have while using NetSight Element Manager.

Using On-line Help

You can use the **Help** buttons in the NetSight Element Manager windows to obtain information specific to the currently active window. When you click on a **Help** button, a window opens which contains information that will assist you in the use of the windows and their associated command and menu options. If a **Help** button is grayed out, on-line help has not yet been implemented for the associated window. From the **Help** menu accessed from the Device View menu bar, you can access on-line Help specific to the Device View (**Help —> Chassis Manager Help** or **Help —> Help Topics**).

Accessing On-line Documentation

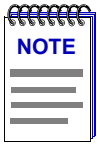

All of the online help windows use the standard Microsoft Windows help facility. If you are unfamiliar with this feature of Windows, you can select ***Help —>How to Use Help*** *from the primary NetSight Element Manager window.*

The complete suite of documents available for NetSight Element Manager can be accessed via a menu option from the primary window menu bar: **Help —> Online Documents.** If you chose to install the documentation when you installed NetSight Element Manager, selecting this option will launch Adobe's Acrobat Reader and a menu file which provides links to all other available documents.

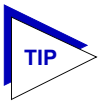

If you have not yet installed the documentation, the ***Online Documents*** *option will not be able to access the menu file. In order to activate this option, you must run the* ***setup.exe*** *again to install the documentation component. See your* **Installation Guide** *for details.*

Documentation Feedback

Send your questions, comments and suggestions regarding NetSight documentation to NetSight Technical Communications via the following e-mail address:

Netsight_docs@enterasys.com

Getting Technical Support

Online Services on the World Wide Web

To locate product-specific information, refer to the Enterasys Web page at the following address:

http://www.enterasys.com

Global Technical Assistance Center

If you have additional questions, contact the Global Technical Assistance Center using one of these methods:

Telephone (24 hours a day, 365 days a year): (603) 332-9400

Fax: (603)337-3075

Electronic Mail: support@enterasys.com

Mailing Address:
Enterasys Networks, Inc.
Technical Support
35 Industrial Way
Rochester, NH 03867

FTP:
ftp.cabletron.com
Login: anonymous
Password: your email address

Chapter 2

Overview of NetSight Element Manager

Key NetSight Element Manager features; customizing the NetSight Element Manager platform; backing up data files

NetSight Element Manager for Windows™ provides a powerful and easy-to-use tool for monitoring and managing Enterasys and Cabletron devices from a Microsoft Windows workstation. Its open architecture and broad feature set also make it an attractive platform for third party development of management applications.

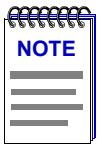

NetSight Element Manager support on the Solaris platform is limited to the installation of NetSight Chassis Manager, either alone or integrated with Hewlett Packard's HP OpenView Network Node Manager, and only for Solaris 2.7. See your **NetSight Element Manager Installation Guide** *for more information.*

About NetSight Element Manager

NetSight Element Manager -- like its predecessor SPECTRUM Element Manager -- is comprised of several independent but interrelated applications:

- The **Network Status** application is responsible for polling nodes in the central database. It provides you with feedback regarding the status of devices on your network through the central NetSight Element Manager user interface via the List Views, Tree Views, and Maps. Node status is propagated right up through these views and onto your desktop via NetSight Element Manager's minimized icon. Key features of this user interface are highlighted in the **Management and Monitoring Workspace**, beginning on page 2-5.
- The **Event Manager** application is responsible for alarm and event management. It receives traps from network nodes and events from other NetSight Element Manager processes, maintains the Log Views in which you view received events, and carries out

any actions that you have configured on reception of a particular trap or event. Alarm and event management is highlighted in **Alarm and Event Handling**, beginning on page 2-11.

- The **Scheduler** tool allows you to schedule unattended data back-ups and network discovery sessions that will launch automatically at a scheduled time and date. With the Scheduler, you can also configure and activate groups of "pager calendars" so that network technicians can be automatically paged in response to network alarms or events. The Scheduler is highlighted in **Scheduler Tool**, beginning on page 2-12.
- The **Remote Administration Tools** suite let you easily perform management tasks (such as TFTP downloads, trap table configuration, or setting up MIB statistics polling or MIB downloads) on an individual device or groups of devices. A MAC Address Locator Tool also allows you to search one or more devices for the port through which a selected MAC address is communicating. The **Remote Administration Tools** section, page 2-12, briefly describes this powerful tool suite.

Among its features are:

- A primary window where you can navigate between the Map View, List View, Tree View, and Log View workspaces to activate windows to monitor and manage your network (such as Map windows and alarm Log View windows). A central menu bar provides access to most management options.
- A graphical interface that lets you sort, rearrange, and resize column displays, select icons for nodes in Map View, List View, and Tree View displays, and customize the toolbar.
- Full-featured on-line assistance, including Wizards for common tasks, context-sensitive help with a glossary and index, and a status bar with hint text. The on-line documentation set is also launchable from the NetSight Element Manager primary window.
- A dockable Workspace area which you can use to access and organize your List View, Map View, Tree View, and Log View windows (as shown in Figure 2-1). You can create folders within a view hierarchy to further nest the contents of your view windows.
- Toolbars which can be docked to the primary NetSight Element Manager window, or undocked as floating tool palettes (as shown in Figure 2-1). Toolbars are automatically activated when a related view is active; you can also choose which toolbars you want to display via the **View—>Toolbars** menu option.

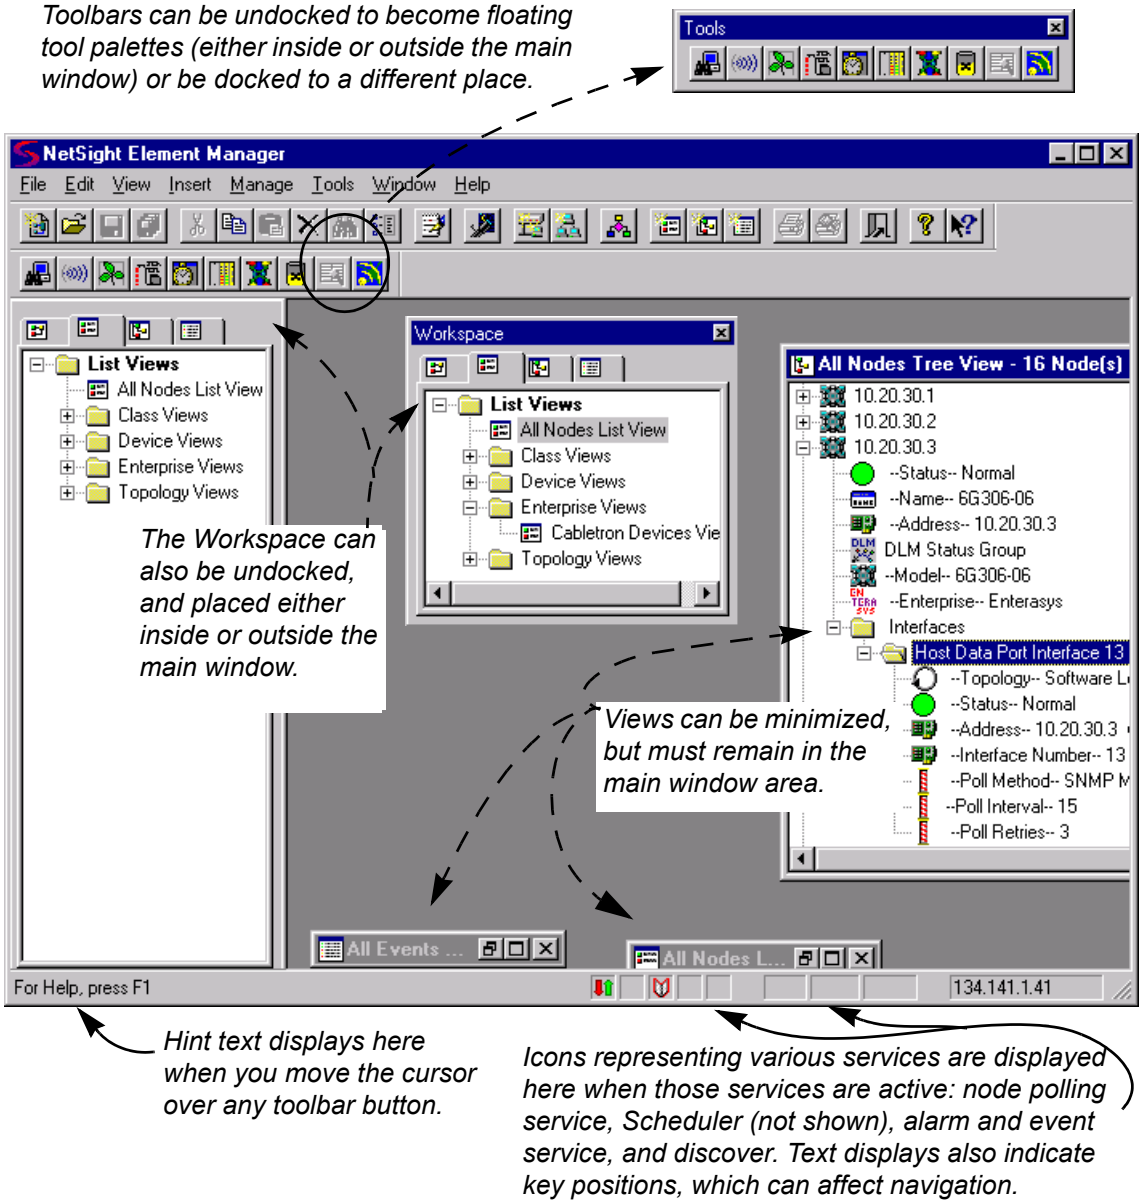
Toolbars can be undocked to become floating tool palettes (either inside or outside the main window) or be docked to a different place.
Tools
NetSight Element Manager
File Edit View Insert Manage Tools Window Help
List Views
All Nodes List View
Class Views
Device Views
Enterprise Views
Topology Views
Workspace
List Views
All Nodes List View
Class Views
Device Views
Enterprise Views
Cabletron Devices Vie
Topology Views
All Nodes Tree View - 16 Node(s)
10.20.30.1
10.20.30.2
10.20.30.3
--Status-- Normal
--Name-- 6G306-06
--Address-- 10.20.30.3
DLM Status Group
--Model-- 6G306-06
--Enterprise-- Enterasys
Interfaces
Host Data Port Interface 13
--Topology-- Software L
--Status-- Normal
--Address-- 10.20.30.3
--Interface Number-- 13
--Poll Method-- SNMP M
--Poll Interval-- 15
--Poll Retries-- 3
The Workspace can also be undocked, and placed either inside or outside the main window.
Views can be minimized, but must remain in the main window area.
All Events ...
All Nodes L...
For Help, press F1
134.141.1.41
Hint text displays here when you move the cursor over any toolbar button.
Icons representing various services are displayed here when those services are active: node polling service, Scheduler (not shown), alarm and event service, and discover. Text displays also indicate key positions, which can affect navigation.

Figure 2-1. The Primary NetSight Element Manager Window, with Docked and Undocked Workspace and Toolbars

- Drag-and-drop functionality for copy or link operations, so that it's easy to transfer data from one management window to the next within the same application (e.g., from a List View to a Map View), or between applications (e.g., from a NetSight Element Manager List View to a device group in Remote Administration Tools).

- Right-mouse button functionality for quick access to commonly-used menu commands (as shown in Figure 2-2).

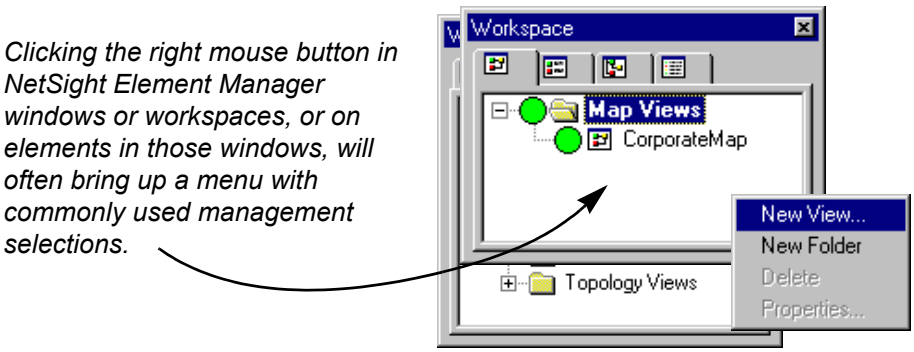
Clicking the right mouse button in NetSight Element Manager windows or workspaces, or on elements in those windows, will often bring up a menu with commonly used management selections.
Workspace
Map Views
CorporateMap
Topology Views
New View...
New Folder
Delete
Properties...

Figure 2-2. Sample Right Mouse Functionality

- Print spooling to the Microsoft Windows Print Manager for record keeping purposes. Printable information includes maps, device lists, alarm statistics, and other statistical information (including RMON data). The Print Preview option is available to let you view the material as it would appear on the page before it is actually spooled to the system printer.
- OLE functionality to embed and link data with management Map files.

The following sections provide an overview of some key NetSight Element Manager functions.

Network Discovery

The Discover application provides several key features; among these are:

Discover Script Capability

By creating Discover scripts, you can store the parameters of a discovery from one management session to the next, so that you can repeat a configured discovery session at any time.

Subnet Discovery

The Discover application is capable of searching based on a subnet address, rather than requiring that a range of IP addresses be specified (although you can enter a range if desired). You can also configure multiple subnets or IP address ranges to be discovered within the same script.

Scheduled Discovery

By linking Discover scripts with the Scheduler application, you can automate the Discover application so that it will run unattended at periodic intervals.

IP or Host Locator

An additional tool linked to but run separately from the Discover application allows you to resolve host names to IP addresses, and vice versa. You can also discover the IP addresses and host names of all devices located on a subnet, and locate both the IP and hostname for a device by its MAC address.

Automatic Creation of the Node Database

Your discovered nodes are automatically entered into NetSight Element Manager's central node database, and the Discover window remains open so that you can view the list of discovered devices. The database is also automatically saved after each discovery.

Secondary Discover and Discover Reset

The first time you run a device discover script, all IP addresses in the subnet or address range you have specified will be queried, and all discovered devices will be listed in the discover window and automatically added to the central node database. This is called a NEW Discovery. The *second* time you run any device discover, however — even a device script which has not yet been run —NetSight Element Manager will perform a SECONDARY Discover, querying only those IP addresses which did not respond to the first discover. This process allows you to use a discover script to periodically survey a subnet, domain, or range of IP addresses for any nodes that might have been down or simply not present during the initial discovery, without the time and bandwidth cost of continually re-polling for nodes which have already been successfully discovered. You can force a return to NEW Discovery mode by clicking the **Reset** button in the Discover Manager window; deleting any node from the central database will also reset the device script counters to NEW Discovery mode.

Discover Wizard

A Wizard is provided to walk you through each step of configuring and performing a discovery.

The Discover application is covered in Chapter 3, **Discovering Nodes**.

Management and Monitoring Workspace

NetSight Element Manager provides four work areas that allow you to choose how to view information about the devices in your central node database: the List View, Tree View, Map View, and Log View. Each provides a slightly different (and user-configurable) look at the central node database, each comes with default views to get you started, and each provides you with the ability to create one or more views specifically designed to display the node information you need to successfully manage your network.

The following sections highlight each of these workspaces, and provide sample illustrations of the workspaces and the windows available from them.

List Views

The List View workspace allows you to view and create specialized lists of the devices in your central node database. A default set of views is provided for you; you can also create your own views, selecting both the nodes and the information about each node you wish to display. Figure 2-3 illustrates a sample List View workspace and an associated List View window.

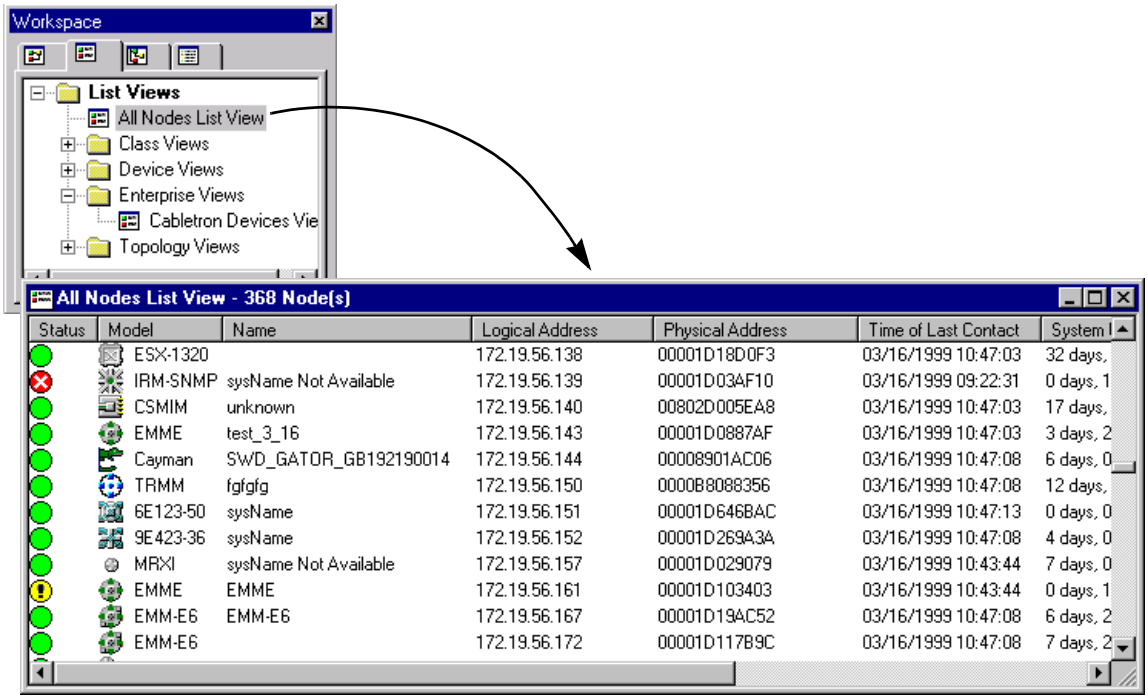

Figure 2-3. Sample List View Workspace and Window

List View features include:

Multiple List Views

You can create multiple List Views to view portions of your node database. The read-only All Nodes List View displays every viewable node in your database, and a set of pre-defined List Views displays various subsets of that database; you can customize this feature by editing a pre-defined view and/or creating an entirely new one to display node information in the way that works best for you. Each List View also provides an Excluded Nodes feature, which allows you to remove one or more nodes from a view without removing them from the database or changing the parameters of the view's filter.

The List View Workspace

Individual List Views are represented as icons within the List View workspace; these icons are arranged in a standard tree hierarchy (using a standard Microsoft Windows tree view control). Double-clicking on a List View icon within the tree opens its associated List View window.

The List View workspace also provides organizational structure to your List Views. The tree starts at the top-level All Nodes List View, and branches into more specifically defined List Views. You can also nest List View icons within folders to further stratify the tree hierarchy.

Detailed Node Information

NetSight Element Manager can display the following information for each node: Status, Model, Name, Label (descriptive alias), Logical (IP or IPX) Address, Physical (MAC) Address, Time of Last Contact, System Up Time, Status Group (DLM – Distributed LAN Management – information), Classes (functional class supported by the node, e.g., router, switch, ATM, etc.), Enterprise (vendor), Topologies supported by the node (e.g., token ring, FDDI, Ethernet, etc.), and Events (whether a trap or event concerning the node has been detected). You can create List Views which select nodes to be displayed based on one or more of these parameters, including subnet, enterprise, class, topology, and status group.

Customizable Node Properties

Each node icon has its own Properties window, which you can use to provide additional descriptive notes for reference, assign a label, change the community name and the interval/retry parameters to be used for polling, and choose icons to represent the node's topology, model, and enterprise.

Tree Views

The Tree View workspace provides access to and allows you to configure a different look at the nodes in your database: the Tree Views. Tree Views display summary information about nodes and their interfaces within a tree hierarchy (using a standard Microsoft Windows tree view control). Figure 2-4 illustrates a sample Tree View workspace and an associated Tree View window.

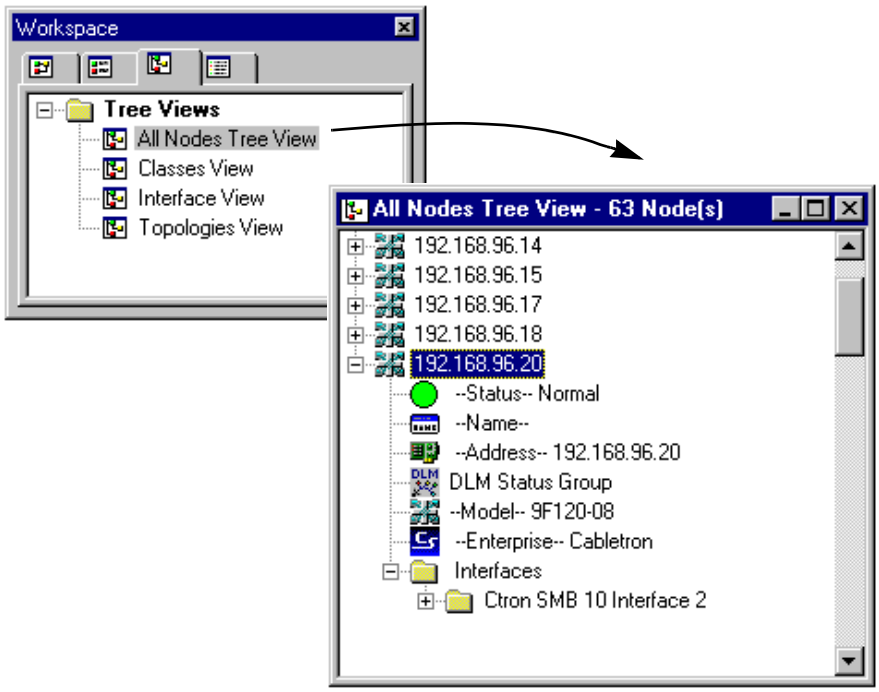

Figure 2-4. Sample Tree View Workspace and Associated Tree View Window

Each node icon in a Tree View window is presented at the top layer in the tree hierarchy. By expanding the display of each node icon, you can view progressively more detailed information about the node (such as polling status, name, IP address, DLM status information, model type, and enterprise) and the interface through which NetSight Element Manager is communicating with the node (such as interface topology, address, index, and polling information).

You can use a Tree View window to glean basic information about each node and its polled interface without bringing up its management window; or, by double-clicking on a node icon, you can open the management window for the selected node.

As with the List View windows, NetSight Element Manager provides a read-only All Nodes Tree View which displays every node in the database; a set of predefined Tree Views (based on classes, interfaces, and topologies) is also provided. The pre-defined views can be edited to suit your needs, or you can create entirely new views as necessary. Also as with List Views, an Excluded Nodes feature allows you to remove one or more nodes from a view without removing them from the database or changing the parameters of the view's filter.

Map Views

Map View windows in NetSight Element Manager provide a third — and perhaps the most flexible — means for viewing information about your network. Using maps, you can create a visual representation of your network: grouping devices by location rather than function, and using a variety of graphical tools (including bitmap backgrounds and drawn objects) to represent connections and other pertinent information about the physical structure of your network. Maps also support Microsoft's OLE (Object Linking and Embedding), which can be used to link or embed data created via other applications into a map document.

NetSight Element Manager's map feature allows you to create a three-dimensional representation of your network by allowing you to show not only relationships between devices, but relationships between *groups* of devices. Individual Map Views can be linked to other Map Views to create a hierarchical structure which is graphically represented in the Map View workspace, again via the Microsoft standard tree control. The overall status of each individual map is also displayed in the workspace; for related maps, each map's status includes the status of any submaps with which it is associated. These status reports are propagated up the hierarchy, whether individual maps are closed or open — giving you a bird's-eye view of your network status at all times.

As with the other workspaces, you can create folders in which to nest groups of Map icons; double-clicking on a Map icon will open the associated Map window. Figure 2-5 illustrates a sample Map View workspace and an associated Map window.

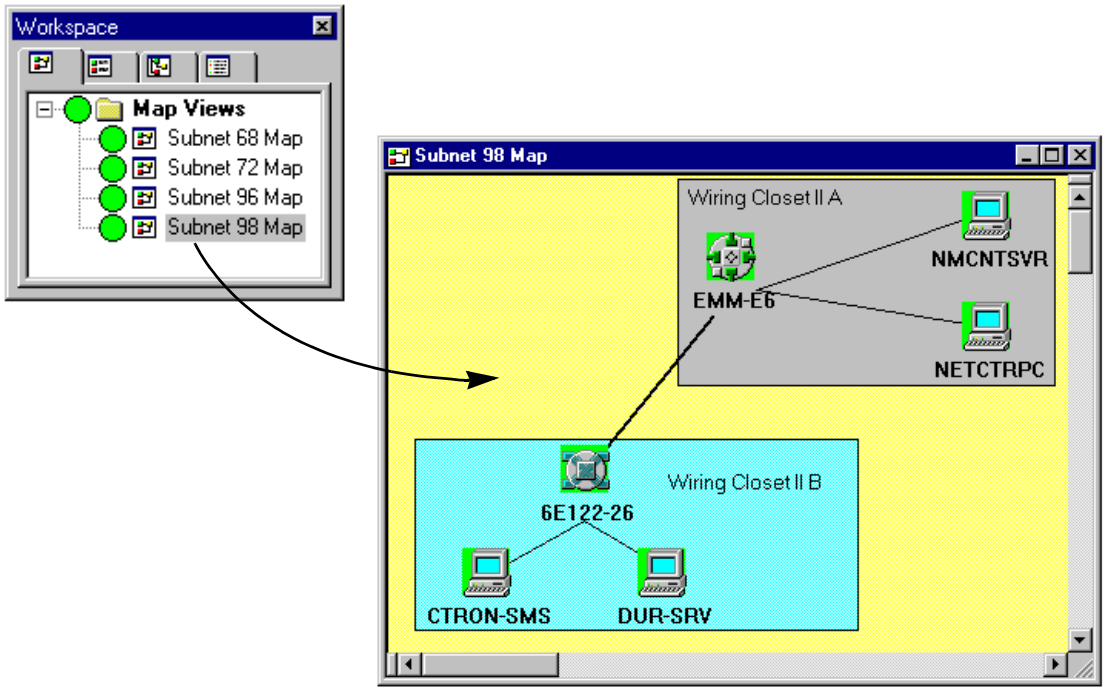

Figure 2-5. Sample Map View Workspace and Map Window

Log Views

NetSight Element Manager provides a Log View workspace, which contains information about traps and software events. Like the other workspaces, the Log View workspace — illustrated in Figure 2-6 — allows you to create, access, and organize the Log View windows which provide the foundation of alarm and event management. Individual Log View windows appear as icons within the Log View tree hierarchy; you can create folders in which to nest individual Log View icons, and double-clicking on a Log View icon will open the associated Log View window. For more on alarm and event management, refer to **Alarm and Event Handling**, page 2-11.

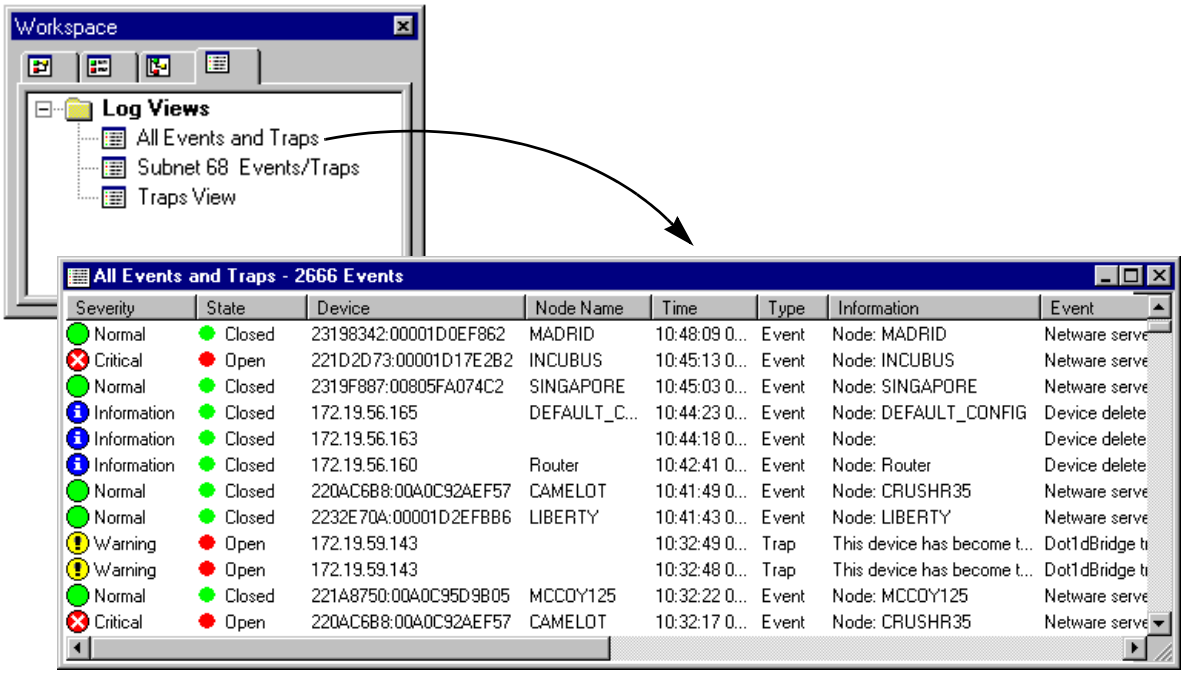
Workspace
Log Views
All Events and Traps
Subnet 68 Events/Traps
Traps View
All Events and Traps - 2666 Events
Severity State Device Node Name Time Type Information Event
Normal Closed 23198342:00001D0EF862 MADRID 10:48:09 0... Event Node: MADRID Netware serve
Critical Open 221D2D73:00001D17E2B2 INCUBUS 10:45:13 0... Event Node: INCUBUS Netware serve
Normal Closed 2319F887:00805FA074C2 SINGAPORE 10:45:03 0... Event Node: SINGAPORE Netware serve
Information Closed 172.19.56.165 DEFAULT_C... 10:44:23 0... Event Node: DEFAULT_CONFIG Device delete
Information Closed 172.19.56.163 10:44:18 0... Event Node: Device delete
Information Closed 172.19.56.160 Router 10:42:41 0... Event Node: Router Device delete
Normal Closed 220AC6B8:00A0C92AEF57 CAMELOT 10:41:49 0... Event Node: CRUSHR35 Netware serve
Normal Closed 2232E70A:00001D2EFBB6 LIBERTY 10:41:43 0... Event Node: LIBERTY Netware serve
Warning Open 172.19.59.143 10:32:49 0... Trap This device has become t... Dot1dBridge t
Warning Open 172.19.59.143 10:32:48 0... Trap This device has become t... Dot1dBridge t
Normal Closed 221A8750:00A0C95D9B05 MCCOY125 10:32:22 0... Event Node: MCCOY125 Netware serve
Critical Open 220AC6B8:00A0C92AEF57 CAMELOT 10:32:17 0... Event Node: CRUSHR35 Netware serve

Figure 2-6. Sample Log Views Workspace and Associated Log View

Excluding Nodes

In addition to providing almost unlimited means for viewing the contents of the central node database, NetSight Element Manager also provides a means for temporarily excluding certain nodes from view — either from a specific view window, or from the central node database itself.

When you exclude a node from a specific view, the node is placed into that view's individual excluded nodes list. It will no longer be visible from that view, but it remains in the database and will be displayed in any other view as appropriate.

When you exclude a node from the central node database, that node is no longer displayed in any view, and is moved to the Excluded Nodes database. Though these nodes can no longer be displayed or managed, they are protected from re-discovery, and they can be re-included in the central node database at any time. Excluded nodes can also be deleted entirely.

The Excluded Nodes database can be viewed via the **View—>Excluded Nodes** menu option available via the primary NetSight Element Manager window; for more information on this feature, see Chapter 4, **List Views**.

Alarm and Event Handling

NetSight Element Manager provides powerful alarm and event handling capabilities via the independent Alarm and Event Service. The Event User Connection Service receives all incoming messages issued by devices configured to send traps to your workstation, as well as system events issued by the NetSight Element Manager Network Status process or another of its component processes (e.g., Discover or Scheduler). Because it is an independent process, the Event User Connection Service can operate without NetSight Element Manager running, receiving (and, where appropriate, acting on) traps issued by your networking devices or events issued by one of NetSight Element Manager's component services.

You can create multiple Log Views to view a single trap/event database; you can also define a separate set of View Filters for each Log View window. The Alarm and Event handler allows you to define two types of event handling actions: foreground actions, which you can apply interactively to a received trap or event; and background actions, which provide unattended checking for a specified trap or event's occurrence, and then activate the pre-determined event handling response.

Trap and event data can be exported to HTML for report generation capability on an internal Web server, and used in conjunction with the Pager alarm notification tool (part of the Scheduler tool described on page 2-12).

For more information about alarm and event handling, see the ***Alarm and Event Handling User's Guide***; for more information about the Scheduler application, see the ***Tools Guide***.

The NetSight Element Manager Tool Suite

NetSight Element Manager includes a number of powerful utility applications — or tools — which ease the task of monitoring and managing your network. Among these tools are Remote Administration Tools, the Scheduler, and MIB Tools.

Remote Administration Tools

Remote Administration Tools is a utility suite which allows you to conveniently perform routine management tasks on a single device — or on a group of devices — from your remote workstation, rather than requiring local terminal access to each device. Among these tasks are:

- Configuring TFTP and BOOTP sessions with a selected device or device group.
- Adding workstation IP addresses to the trap table(s) of a selected device or device group.
- Searching a group of devices for the port through which a specified MAC address is communicating.
- Performing MIB SETs, using a template of MIB objects and values, on one or more devices.
- Performing MIB GETs, using a template of MIB objects, from one or more devices for statistics and trend analysis.
- Performing a soft reboot on selected devices.
- Changing access community names on selected devices.

For more information about these tools, see the ***Remote Administration Tools User's Guide***, included in your documentation set.

Scheduler Tool

The Scheduler tool allows you to schedule certain administrative actions so that they can be performed while your workstation is unattended. These actions include:

- Performing discovery sessions (using scripts defined using the Discover application).
- Launching selected programs, including TFTP downloads, MIB statistics recordings, and MIB loads via the Remote Administration Tools application.
- Backing up files (NetSight Element Manager data files, or any other files you choose).
- Setting up automated paging (by defining individuals or groups of individuals to be paged, and specifying their availability to be paged). The paging feature is integrated with NetSight Element Manager alarm and event handling, so that a scheduled page can be performed as part of an interactive foreground action, or an unattended background action.

For more information about the Scheduler tool, see the ***Tools Guide***, included in your documentation set.

MIB Tools

The MIB Tools suite of MIB browsing and editing tools comprises:

- The **MIB Browser** — a simple graphical interface you can use to contact and query SNMP nodes on your network and examine their supported MIBs (as long as the MIBs are imported into the MIB Tools database). MIB Browser also lets you conveniently perform SETs on contacted devices.
- **MIB Details** — a display of the descriptive information for management objects that you query from a node.
- The **MIB Editor**, which lets you manipulate your database of network vendor's MIBs. An extensive MIB database is provided with MIB Tools; you can input new MIBs (stored in ASCII text format) into this database, or you can delete any unnecessary MIBs (or MIB portions) so that the database maintains the precise MIB information that you need to manage your nodes.
- The **Device Manager**, which allows you to maintain a database of node IP addresses that you frequently contact, and displays summary information about these nodes.
- A **Preferences** option, which lets you view and change configuration settings for the MIB Tools suite.

Customizing NetSight Element Manager

There are a number of ways you can adjust NetSight Element Manager's settings to customize its operation in the way that works best in your management environment. These are detailed in the following sections.

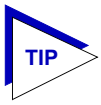

*You can also do some batch customization outside of the NetSight Element Manager platform; see Appendix A, **Batch Customization via AddTool, AddImage, and AddDevType**, for more information.*

Adjusting Default Settings

The Options window controls the default settings for NetSight Element Manager and some of its components. These settings control the operation of many basic features and functions in NetSight Element Manager, and they are referenced throughout the document set. You may want to adjust these settings so that they better suit your specific network management needs.

The Options window — accessed by selecting **Tools—>Options** from the primary window menu bar — has seven tabbed pages, each of which controls a set of related functions in NetSight Element Manager. Each tabbed page and the options it controls is discussed in the following sections.

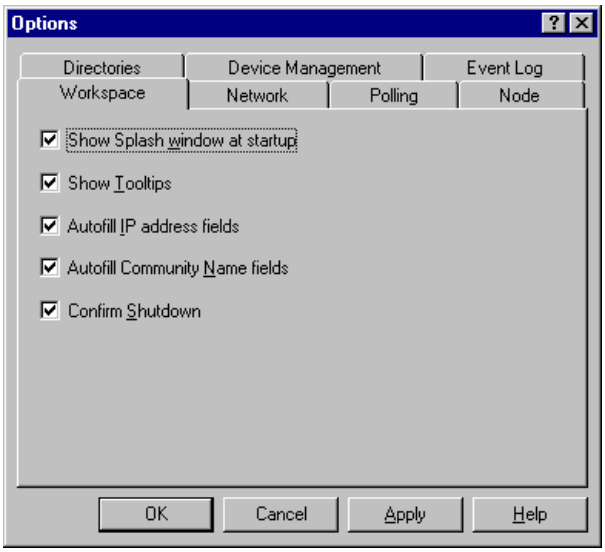

Figure 2-7. The Options Window (with Workspace Options Selected)

Workspace Options

The Workspace options (displayed by default, as illustrated in Figure 2-7) control various settings for the NetSight Element Manager interface. You select or deselect each option by clicking the associated checkbox. A checked box indicates that the associated option is selected; a blank box indicates that option is deselected.

Show Splash window at startup

The Splash window is an identifying screen that will briefly appear when you initialize NetSight Element Manager. Selecting this option will cause the screen to be displayed; deselecting it will suppress the screen.

Show Tooltips

Tooltips are small pop-up windows that provide user assistance by identifying controls such as toolbar buttons. Selecting **Show Tooltips** will cause these to be displayed when the mouse cursor is dragged over a toolbar icon; deselecting the option will suppress them. You may want to leave this option selected until you are comfortable visually identifying all the toolbar options by their icons.

Autofill IP address fields

When this option is selected, the dotted decimal IP address octets set via the Network options (or via the Welcome Wizard) will be automatically entered in any field that requires an IP address.

Autofill Community Name fields

When this option is selected, the community name text set via the Network options (or via the Welcome Wizard) will be automatically entered in any field that requires a community name.

Confirm Shutdown

By default, when you shut down NetSight Element Manager (via the **File—>Exit** menu option, the **Exit** toolbar icon, or the Close button in the main application window), a message window opens asking you to confirm your action with a **Yes** (exit) or **No** (remain open). Deselecting this option will cause the confirmation window to be suppressed and the application will immediately shut down; leaving it selected will cause the confirmation window to be displayed.

Network Options

The Network options (Figure 2-8) allow you to set the default Community Name, IP Address prefix, and Address Mask that will be used when you attempt to add entries into the node database (either manually or through a discovery).

TIP

When you initialize NetSight Element Manager for the first time, a Welcome Wizard will guide you through setting up the product. The information entered through the Welcome Wizard appears (and can be edited) here.

Community Name

The Community Name is a basic security mechanism implemented in SNMP. It acts as both a local and remote password for IP devices, and provides different levels of access (i.e., read only, read-write, and super user) to the device's information. Community names and their level of access must be configured at each device. Generally, the default community names for a device are set at "public" for all levels of access.

TIP

If you have adequate privileges, there are several ways you can edit a device's community name: via Local Management, as described in your local management documentation; via any MIB utility — like NetSight Element Manager's MIB Tools suite — that allows you to perform SNMP GETs and SETs; or via the Community Names component of the Remote Administration Tools.

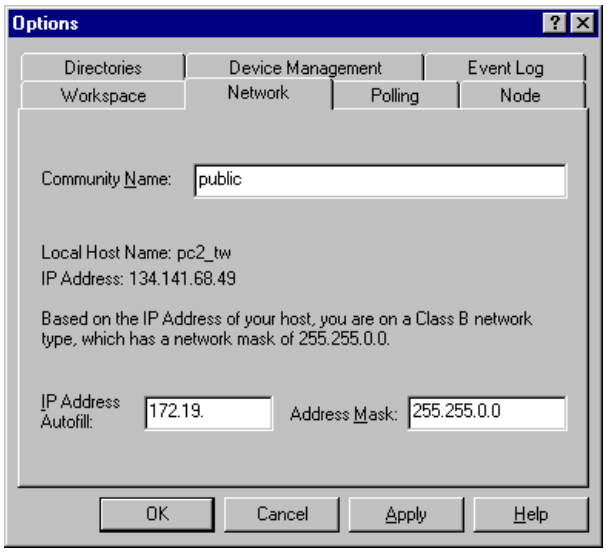

Figure 2-8. The Network Options

NetSight Element Manager uses the Community Name to gain access to devices that you want to manage. You specify the community name you want to use to access a device when the device is inserted (as documented in Chapter 4, **List Views**) or discovered (as documented in Chapter 3, **Discovering Nodes**). You can also change the community name NetSight Element Manager uses to communicate with a device via the Device Properties windows (also documented in Chapter 4, **List Views**). For full device management, you should select a community name which provides super user access.

The community name set in the Network Options window serves as the default value entered in any field which requires a community name; however, you can always enter a specific community name on a case-by-case basis in any of the windows that require one.

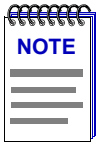

To check whether a device is active on the network — and consequently to turn the device icon green on Map Views, List Views, and Tree Views — NetSight Element Manager issues polls for device interface information. This type of poll does not require community name access to the device, since it is merely a check of network health. If you cannot access a device's management windows — even though its icon appears green on the map — check that the community name you are using to access the device still has valid read, read/write, or super user access.

To set the default community name:

1. Click in the **Community Name** text field, and type in a community name which will be used as NetSight Element Manager's default value.

See your local management documentation or the ***Remote Administration Tools User's Guide*** included with your documentation for more information on community names.

IP Address Autofill

As with community names, an IP address is required to add devices to, or modify existing devices in, the central node database. (The Insert Device and Device Discover Properties windows both have an IP address entry field.) The IP Address Autofill option allows you to specify the network or subnet identifying octets (in dotted decimal IP format) you wish to be entered by default in any required IP Address field (again, you can change the IP address on a case-by-case basis).

To set the default IP Address Autofill prefix:

1. Click in the **IP Address** text field, and type in the network or subnet identifying octets (in dotted decimal IP format) which you wish to use as a prefix.

Address Mask

An address mask is used to determine whether a given destination IP address (in this case, that of managed devices) exists on the same network or subnetwork as a given source IP address (in this case, that of yourNetSight Element Manager workstation). The Address Mask field lets you set a default value which will be used when NetSight Element Manager polls for IP devices during a Network Discovery.

The appropriate address mask depends upon your network configuration. If your network contains subnets, entering the appropriate *subnet mask* in this field will identify the network, subnet, and host portions of the IP addresses in the address range. For example, on a Class B network containing an IP address of 132.177.118.24, a subnet mask of 255.255.255.0 would identify the first two octets (132.177) as the network portion of the address, the third octet (118) as the subnet portion, and the final octet (24) as the host portion.

If your network does not contain subnets, entering the appropriate *network mask* in this field will identify the network portion of the IP addresses in the address range. For example, on a Class B network containing an IP address of 132.177.118.24, a network mask of 255.255.0.0 would identify the first two octets (132.177) as the network portion of the address.

To set the default Address Mask:

1. Click in the **Address Mask** text field, and type in the network or subnet identifying octets which you wish to use as the mask.

Polling Options

The Polling options (Figure 2-9) let you configure how NetSight Element Manager contacts nodes in your database (and consequently how often their contact status is refreshed in List View, Tree View, or Map View windows).

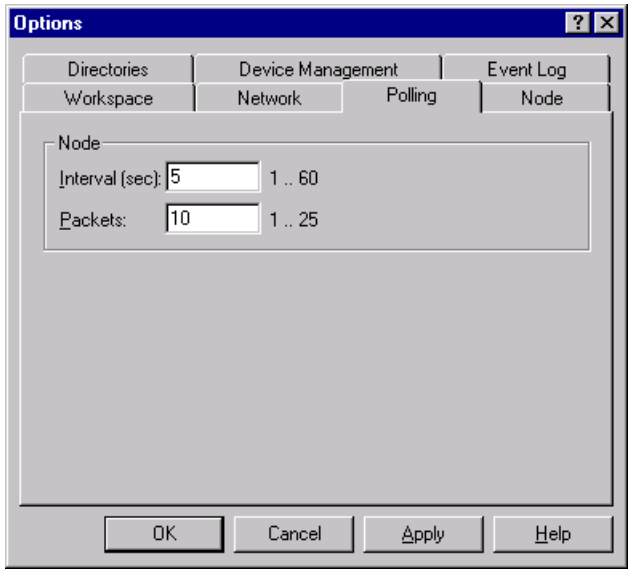

Figure 2-9. The Polling Options

Node Polling

Node polling establishes how often the contact status between NetSight Element Manager and the IP devices in your node database will be refreshed. You configure node polling by setting the interval for a polling cycle and the number of nodes that will be polled during that interval. In essence, you use the node polling options to limit the amount of network traffic that NetSight Element Manager generates as it polls devices for their status.

The above figure shows the defaults for NetSight Element Manager: a 5-second polling cycle with 10 poll packets issued within that cycle. With a database of 100 managed nodes, this means that it will take 50 seconds to poll the entire node database and fully refresh the contact status of nodes in the List Views, Tree Views, and Map Views — since there will be 10 iterations of the 5-second polling cycle to determine the contact status with all 100 nodes.

NOTE

The status of polled nodes is maintained as is between polling cycles. If a node goes down during a polling cycle, its icon will remain Green/Up until NetSight Element Manager re-polls it and determines the loss of contact. For an immediate update of a specific node's status, highlight the device, then select ***Manage—>Query Status*** *from the primary window menu bar, or* ***Query Status*** *from the right mouse menu.*

You can configure polling for an individual node so that it is polled *less* frequently than the refresh time required to poll the entire node database. For example, if one of your 100 managed IP nodes was at the far end of a T1 link, and you would like that node to be polled hourly (rather than at the refresh rate established via Node polling), you could set an individual poll rate for that device to be 3600 seconds (via the Device Properties

window, as described in **Configuring Node Interface Poll Settings** in **Chapter 4**). This would override the polling refresh rate of the database (once per 50 seconds), by instead having NetSight Element Manager issue a once-hourly poll to that individual node.

However, you cannot configure polling so that contact status with an individual node is refreshed *more* frequently than the refresh rate established by the Node database polling settings. Even if you configured one of your 100 managed nodes to be polled every 10 seconds for contact status, it would still be subject to the refresh rate of the entire node database. In actuality, its contact status would only be updated once during the 50 seconds it takes to fully poll the node database.

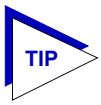

If you have a mission-critical node, you can keep Chassis Manager running to be constantly appraised of its contact status. Because the Chassis Manager application uses its own polling mechanism (independent of the node database polling in NetSight Element Manager), you can determine the status of the node by using the Contact Status symbol in Chassis Manager's Logical View window. Refer to your individual device-specific **User's Guides** *for more information on using Chassis Manager; refer to the* ***Device Management Options*** *section, following, for more information on setting the Chassis Manager poll settings. You can also issue an immediate status poll at any time via the* ***Query Status*** *options described above.*

Use the following fields to establish the refresh rate of polling for all nodes in your node database:

Interval (sec)

The amount of time that elapses between groups of poll packets. The number of intervals required to complete a poll cycle will depend on how many devices are in your node database, and how many are being polled during each interval.

To set the Interval value:

1. Click in the **Interval** text field, and type in the number of seconds you want NetSight Element Manager to wait between each set of polling packets. The range is from 1 to 60 seconds; the default is 5 seconds.

Packets

The number of poll packets that will be sent to devices eligible for polling during each polling interval.

To set the Packets (number of polled devices) value:

1. Click in the **Packets** text field, and type in the number of IP devices you want polled during each poll Interval. The range is from 1 to 25.

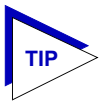

To determine the amount of time it will take NetSight Element Manager to complete one poll of all nodes in your database, divide the number of nodes in your database by the ***Packets*** *value and multiply by the* ***Interval (sec)*** *value.*

Node Options

The Node options (shown in Figure 2-10) let you configure the default contact poll settings that will be established when IP devices are added to the node database.

Default Poll Settings

Unlike the poll settings configured via the Polling Options tabbed page, the Node Options default poll settings determine the default poll values that will be assigned to each individual node via its Properties window. In most instances, these default poll settings will have no bearing on how often the node is actually polled, since the **Node Polling** settings determine the amount of time it takes NetSight Element Manager to poll every device in the database.

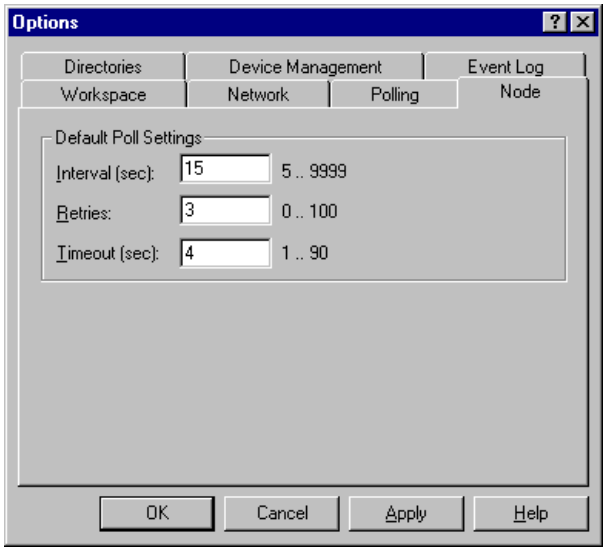

Figure 2-10. The Node Options

In general, it is advisable to leave the Default Poll Settings at their initial values; if you have a few specific devices you'd like to poll *less* frequently than the configured Polling Options values would achieve, use the individual nodes' Properties windows to edit these default settings.

The Default Poll Settings let you configure the following:

Interval (sec)

The default number of seconds configured for polling an individual IP device. Note that in general, the polling settings for updating the status of individual IP devices will be overridden by the Node options configured in the Polling page. The interval set in this field will only take precedence if it is *longer* than the duration required to poll the entire node database (as configured via the **Polling Options** page).

To change the interval (if desired):

1. Click in the **Interval (sec)** text field, and type in the poll interval to be used as a default when adding new devices. The range is from 5 to 9,999 seconds.

Retries

If individual device polling is in effect, this will determine the default number of times after a failed poll that NetSight Element Manager will re-poll a device before declaring it in a Critical (contact unavailable) state. To change the default Retries setting:

1. Click in the **Retries** text field, and type in the number of polls which will be issued to an individual device before declaring a contact unavailable status. The range is from 0 to 100 retries; the default is 3 retries.

Timeout (sec)

This field determines the default duration between retry polls to an individual device. To change the Timeout value:

1. Click in the **Timeout (sec)** text field, and type in the duration (in seconds) between re-polls after the initial poll failure. The range is from 1 to 90 seconds; the default is 4 seconds.

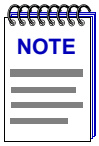

This Timeout value can only be configured here, for all nodes; individual Node Properties windows do not provide the option of setting this value.

Directories Options

The Directories page lists the default locations used by NetSight Element Manager and its components to store and retrieve files. It also allows you to change the default Map and Export directories.

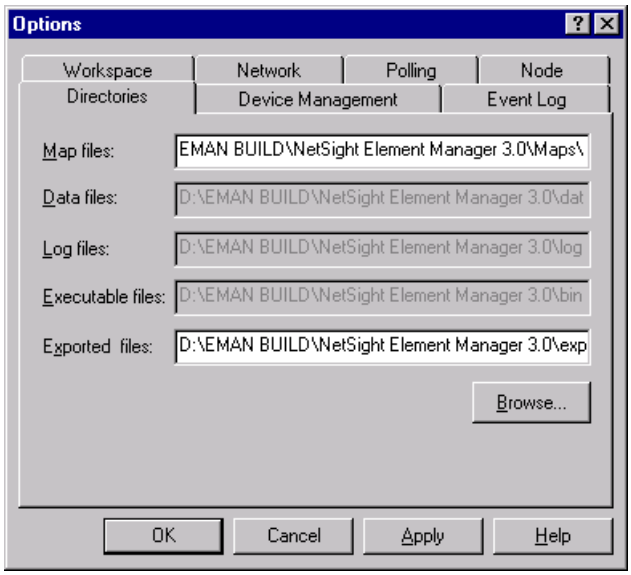

Figure 2-11. The Directories Options

Map files

This field specifies the default directory that is used when you save a map. You can either override the default value in the Save Map window as each new map is saved, or you can set a new default location here. Refer to Chapter 6, **Creating Network Maps**, for more information on creating and saving maps. The default directory is \NetSight Element Manager x.x\Maps.

To specify the default Map directory:

1. In the **Map files** text field, type in the full directory path for the directory in which you want to save map files by default. Be sure the directory already exists, or the change will not be accepted.

   or

   Click **Browse**. A standard Microsoft Browse window opens. Use the directory tree to select the directory (folder) which you would like to use as the default.

TIP

If you change the default map location, map files stored there will still be automatically backed up with the Data File Manager; however, if you also have map files in the original directory (or in any other directory), those will not be automatically backed up. If you store map files in more than one directory, be sure to back them up manually.

Data files

This read-only field indicates the directory where NetSight Element Manager stores its workspace configuration information (including information about the nodes in the List Views and Tree Views). The default directory is \NetSight Element Manager x.x\Data.

Log files

This read-only field indicates the directory used to store log files generated by NetSight Element Manager's processes or applications (e.g., the alarm and event handling service). The default directory is \NetSight Element Manager x.x\Log.

Executable files

This read-only field indicates the directory whereNetSight Element Manager stores its executable files and dynamic link libraries (DLLs). The default directory is \NetSight Element Manager x.x\Bin.

Exported files

This field displays the default directory to be used in any export action taken from within NetSight Element Manager or one of its component applications (like RMON), regardless of the export file type (*.csv, *.txt, *.htm, etc.). The default directory is \NetSight Element Manager x.x\Export.

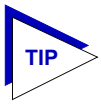

If you change the default export file location, files stored there will still be automatically backed up with the Data File Manager; however, if you also have exported files in the original directory (or in any other directory), those will not be automatically backed up. If you store exported data files in more than one directory, be sure to back them up manually.

Device Management Options

The Device Management options (Figure 2-12) let you configure the poll settings to be used by the Chassis Manager and SmartSwitch/Matrix chassis management applications.

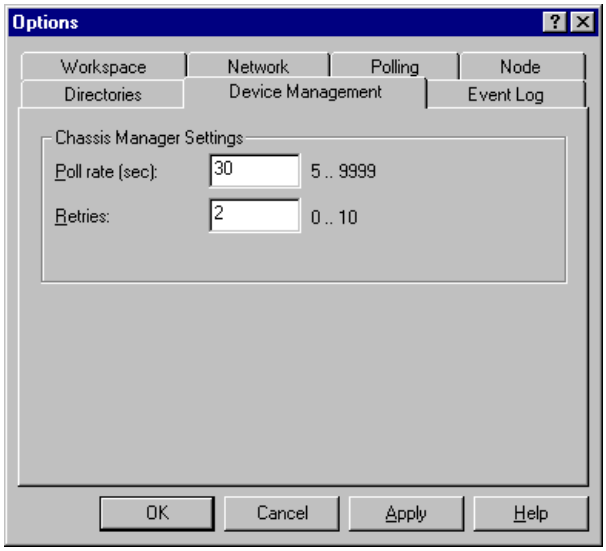

Figure 2-12. The Device Management Options Window

Chassis Manager Settings

The Chassis Manager settings let you establish how the Chassis Manager and SmartSwitch/Matrix Chassis View applications poll a monitored device. There are two settings which you can configure:

Poll Rate

The interval, in seconds, at which the Chassis Manager or SmartSwitch/Matrix Chassis View will poll a device to check whether it is alive on the network (i.e., has a green contact status).

To change the poll rate:

1. Click in the **Poll Rate** text field, and type in the interval to be used when Chassis Manager polls a device for its contact status. The range is from 5 to 9999 seconds; the default is 30 seconds.

Retries

The Retries field determines the number of polls that the Chassis Manager or SmartSwitch/Matrix Chassis View will issue to a device after the first failed poll before declaring a lost contact state. To change the Retries setting:

1. Click in the **Retries** text field, and type in the number of polls which must fail (after the initial poll failure) before Chassis Manager declares that contact is unavailable. The range is from 0 to 10 retries; the default is 2 retries.

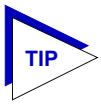

You can also edit the ***ctron.ini*** *file (located in the* /windows *directory for Windows 2000, or in the* /winnt *directory for Windows NT) to specify the interval between retry polls from Chassis Manager or the SmartSwitch Chassis View. Alter the* ***BackgroundTrysInterval=X*** *line, where X is the number of seconds between retry polls. The default setting is 15 seconds.*

Event Log Options

The Event Log options (Figure 2-13) allow you to set the size of the alarm and event service's event cache (which stores trap and event information). The event cache is a memory-mapped data file that is activated when NetSight Element Manager's alarm and event service initializes. All traps and events that are received are uniquely identified and stored in this file. The file is a wrap-around buffer which overwrites the earliest entries with each new incoming trap or event. Refer to your ***Alarm and Event Handling User's Guide***, included in your documentation set, for more information about the event cache and other aspects of alarm and event handling.

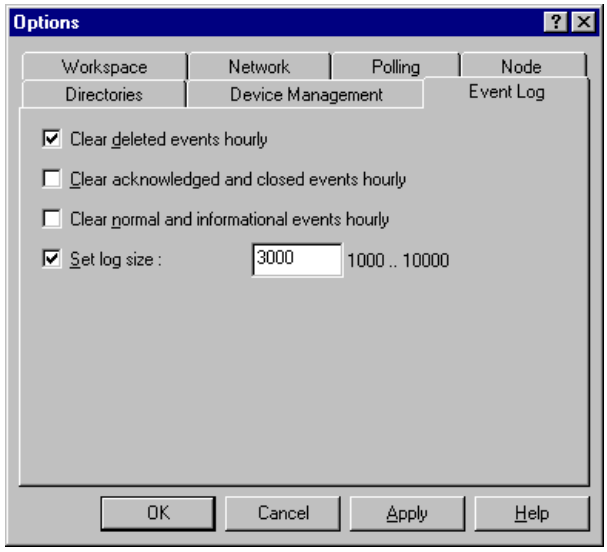

Figure 2-13. The Event Log Options Window

You can establish an hourly purge of the event cache file by specifying that traps or events which meet certain conditions be deleted. These conditions are defined by three Event Log options you select or de-select by clicking in the associated checkbox. The fourth option — **Set log size** — allows you to control the overall size of the cache.

Clear deleted events hourly

Checking this option will activate an hourly purge of all events marked as deleted. If this option is deselected, events marked as deleted will simply be overwritten as the cache wraps at their position.

Clear acknowledged and closed events hourly

Checking this option will activate an hourly purge of all events that have been administratively acknowledged or closed during the previous hour. If this option is deselected, no such purge will take place, and acknowledged and closed events will simply be overwritten in the normal course of the event cache wrap.

Clear normal and informational events hourly

Checking this option will activate an hourly purge of all events received during the previous hour that have a severity level of Normal or Informational. If this option is deselected, no such purge will take place, and normal and informational events will simply be overwritten in the normal course of the event cache wrap.

Set log size

This option will allow you to change the size of the event cache from between 1000 and 10000 entries; the default size of the cache is 3000 entries.

To set the options in the Event Log Options window:

1. To activate the **Clear deleted events**, **Clear acknowledged and closed events**, or **Clear normal and informational events** options, click to check the appropriate selection.
2. To set a new event cache size:
   a. Click to check the **Set log size** option. Note that if this option is not selected, the cache size text field will not be editable.
   b. Click in the associated text field and type in the desired cache size. The range is 1000–10000 entries.
3. Click **Apply** to set your selection and remain in the Options window, or **OK** to set the selection and exit the Options window.

Once you apply any change, the appropriate changes are immediately made to the event cache (e.g., all normal and informational events will be deleted), and from that point the cache will be updated hourly as configured. Note that each time you make and set a change in the Event Log Options window, you will restart the hourly purge cycle based on the time of configuration.

Customizing the Toolbar

NetSight Element Manager provides access to a number of tool applications via its Tools menu and accompanying Tools toolbar. You can customize this access by adding items to or removing items from both the menu and toolbar, or editing their tool tip, hint text, and/or icon image. This process is described below.

*You can also do some batch customization of the Tools menu and toolbar outside of the NetSight Element Managerplatform; see Appendix A, **Batch Customization via AddTool, AddImage, and AddDevType**, for more information.*

To access the customize window:

1. Select **Tools—>Customize** from the primary window menu bar. The Customize window, Figure 2-14, opens.

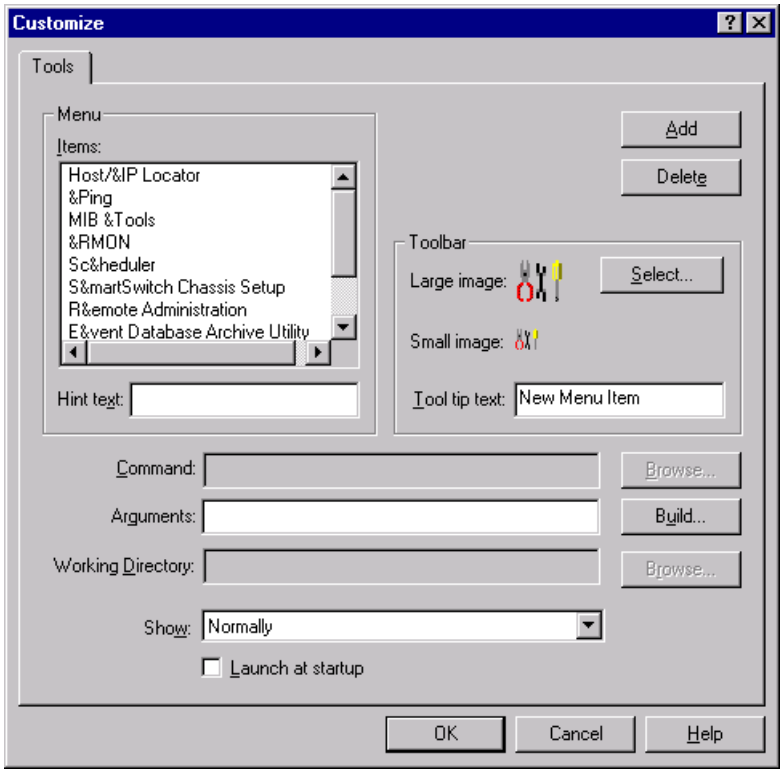

Figure 2-14. The Customize Window

The top portion of the window allows you to define the menu pick, hint, and tool tip text that will label your new selection, as well as the graphic image that will be used as its tool bar icon. The lower portion of the window allows you to assign the appropriate executable and any necessary arguments; it also allows you to define certain launch parameters. You can use these same fields to edit any existing entry in the toolbar or menu.

TIP

Information about the configuration of the Tools menu and toolbar is stored in the Tools.dat file (\NetSight Element Manager x.x\data); to preserve your customization, be sure to back up this file. See ***Backing Up Data Files****, beginning on page 2-30, for more information.*

Adding a Tool Entry

When you add an entry to the Tools menu, you must also add a corresponding icon to the Tools toolbar, and vice versa. To do so:

1. Click on the **Add** button in the upper-right corner of the window. A new menu item with default text (**&New Menu Item**) will be added to the Menu Items list box, with an edit box around it.

2. In the edit box provided, enter the text you want to appear in the Tools menu for the new entry. Place an ampersand (&) before the character you'd like to use for the access key.

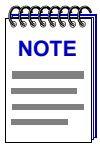

If you select an access key that is already in use, a message window will note this when you try to enter your new item name. If you choose to configure an access key, it must be unique or the new Tools item cannot be added.

3. In the **Hint text** field, enter the text you want displayed in the status bar at the bottom left of the primary NetSight Element Manager window when the menu item is selected or the cursor is positioned over its associated toolbar icon. This information is optional; you can also leave the field blank.
4. In the **Tool Tip text** field, enter the text you want displayed in the tooltip pop-up when the cursor is positioned over the associated toolbar icon. This field is also optional, and can be left blank.
5. To select an image to use for the toolbar icon, click on **Select** to launch the Select Image window, Figure 2-15. Once there:

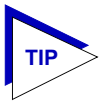

If the executable you have selected already has an icon associated with it, that icon will be displayed by default; you can use the icon image provided, or change this default via the following steps.

   a. Click **Browse** and in the resulting File Open window, select the bitmap file you wish to use as a basis for the toolbar icon, then click **Open**. The selected bitmap image will be displayed in the **Images:** display area in the window. (If the image is too big in either dimension to be fully displayed, a message window opens, indicating that the image will be cropped to fit.)
   b. Click the **Small Image** radio button, then move the resulting square selection cursor around on the image to select the portion you wish to display for the toolbar icon, and click again to make the selection. The image portion you have selected will be displayed below the **Small Image** radio button. Move the cursor and click again to change this selection until you find the one you want.
   c. If you wish, click the **Large Image** radio button, then move the resulting square selection cursor around to select the image portion you wish to display for a large icon, and click again to make the selection. The image portion you have selected will be displayed below the **Large Image** radio button. Move the cursor and click again to change this selection until you find the one you want. This image isn't displayed anywhere except in this window, but you may want to go ahead a select one for consistency. If you don't select one, the default image will remain in place.

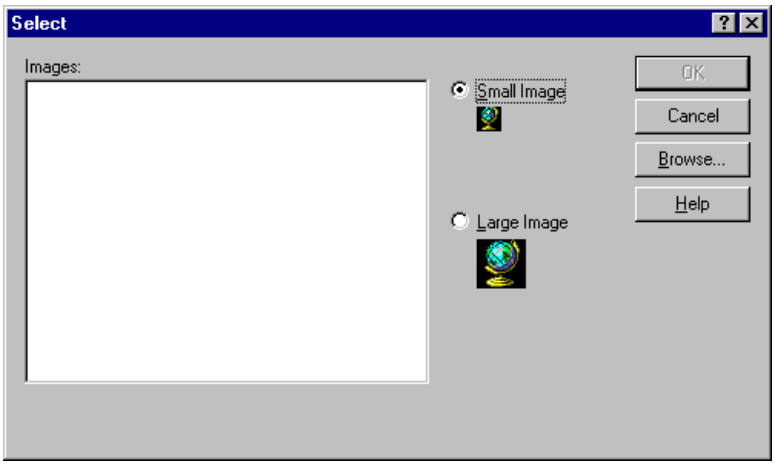

Figure 2-15. Select Image Window

d. Click **OK** to save your image selection(s) and return to the Customize window.

6. Use the **Browse** button to the right of the **Command** field to locate the executable you wish to launch with your new Tools menu selection. The executable path and file name cannot be entered directly into the **Command** field; you must select it via the Browse window.

7. Use the **Arguments** field to add any arguments that can be accepted by your chosen executable and that you want added by default each time the icon or menu pick is selected.

TIP

*Clicking on the **Build** button will launch a Build window that may help you construct the appropriate argument string; for more details on how to use this window, see **Building an Argument List** in Chapter 4, **List Views**.*

*Note that if you select an argument symbol that will extract information from the NetSight Element Manager database (for example, **%i** to insert an IP address, or **%c** to insert a community name), the menu item and tool bar icon for your application will only become available when a device is selected in a List, Tree, or Map View — this will guarantee that the appropriate argument information can be obtained.*

If you select a specific value for a build argument, or if you do not specify any build arguments, your Tools menu item and its accompanying icon will always be available.

8. If you wish, you may also specify a **Working Directory** for your tool. Use the **Browse** button to the right of the **Working Directory** field to locate the appropriate directory. The directory name cannot be entered directly into the field; you must select it via the Browse window.

9. In the **Show** field, use the drop-down list to select the default display mode for your tool application: **Normally**, **Hidden**, **Minimized**, or **Maximized**.
10. Click to select the **Launch at Start-up** option if you want this tool launched automatically with NetSight Element Manager.
11. Click **OK** to add your new option.

Editing a Tool Entry

You can also edit any of the configured values for a Tool entry (either one you have added, or one of the default entries provided); to do so:

1. Open the Customize window by selecting **Tools—>Customize** from the primary NetSight Element Manager window.
2. In the Menu Items list box, click to select the item you wish to edit. The remaining window fields will fill with the values set for the selected item.
3. Edit the existing values as desired; see the previous section for details.
4. Click **OK** to save your changes and exit the Customize window.

Removing a Tool Entry

You can also delete any existing Tool entry, as follows:

1. Open the Customize window by selecting **Tools—>Customize** from the primary NetSight Element Manager window.
2. In the Menu Items list box, click to select the item you wish to remove, then click on **Delete**. The selected item and its associated parameters will be removed from the window.
3. Click **OK** to save your changes and exit the Customize window.

Backing Up Data Files

We strongly recommend that you back up your NetSight Element Manager files on a regular basis. The Scheduler will do that automatically (see the **Scheduler** chapter in the ***Tools Guide*** for more information). This section describes how you can also perform this function manually, with the Data File Manager.

The Data File Manager is preset to backup specified files and file types in five directories located in the \NetSight Element Manager x.x directory: \Data, \Bin, \Filters (and its automatically-created subdirectories), \Maps (or your configured Maps directory), and \Export (or your configured Exported Files directory). An editable configuration file determines the specific files and file types that will be backed up; you can edit this file to add additional files or file types, but they must reside in one of the five directories specified above.

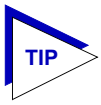

*If you change the default Maps and/or Exported Files directory (via the **Directories Options** available in the **Tools—>Options** window), the Data File Manager will automatically back up any specified files that reside in the new default directories. Any files that reside in the original default directories, however, will no longer be backed up.*

The following files and file types are included in the Backup Configuration file by default; they will be backed up and maintained in your backup directory as long as they reside in one of the defined backup directories.

- **rmflt.mdb** and **rmflt.ldb** — These database files maintain RMON packet capture filter definitions.
- ***.csv**, ***.htm**, ***.mrf**, and ***.txt** — Files which contain exported data
- ***.map** — .MAP map files (viewable format)
- ***.mmp** — SmartSwitch 9000 Chassis Configuration files (from SPECTRUM Element Manager version 2.0)
- ***.dmf** — SmartSwitch and Matrix Chassis Configuration files
- ***.mrm** — .MRM map files (import format)
- ***.filter** — Files with information for viewing filters for log view events
- ***.action** — Files containing definitions for configured foreground actions — actions that must be manually initiated in response to an event
- ***.backgroundaction** — Files containing definitions for configured background actions — actions that are executed automatically in response to specific events
- **ae_log.mdb** — Alarm and Events log database
- **Classes.dat** — Contains definitions for node classes
- **Discovers.dat** — Scripts used in network discovers
- **EvtCache.dat** — The cache which stores the most recent traps and events
- **Exnodes.dat** — Excluded nodes
- **Groups. dat** — Information on DLM status groups
- **Intrface.dat** — Interface information displayed in Tree Views, List Views, and the Node Properties window
- **Nodes.dat** — Node information displayed in Tree Views, List Views, and the Node Properties window
- **Port.dat** — Reserved for future use
- **RemoteAd.mdb** — The Remote Administration Tools database
- **Sessions.dat** — Definitions for Netstat, NetSight Element Manager, and scheduler connections

- **Subnet.dat** — Subnet definitions
- **Tools.dat** — Information on items in the Tools menu
- **Wrkspace.dat** — Configuration of window placement and other appearance properties within the primary NetSight Element Manager window.
- **CSMIBDS.MDB** — MIB database.
- ***.def** and***.trapdef** — Trap definition files.

To access the Data File Manager:

1. Exit any current management sessions. NetSight Element Manager must be closed for the backup utility to operate.
2. Select the **Data File Manager** icon in your NetSight Element Manager program group. The Data File Manager window opens, as illustrated in Figure 2-16.

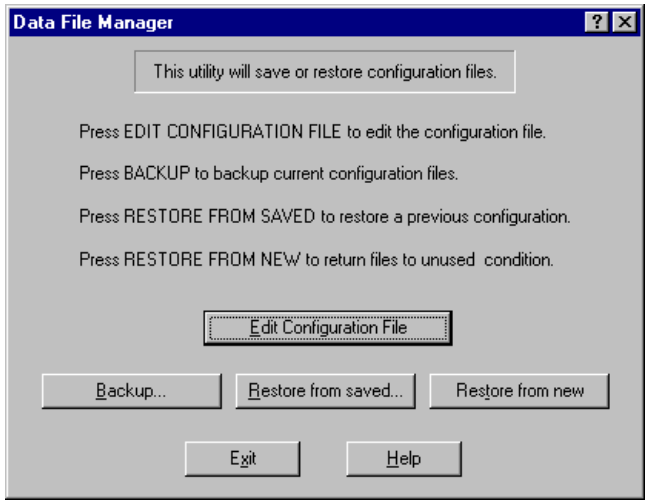
Data File Manager
This utility will save or restore configuration files.
Press EDIT CONFIGURATION FILE to edit the configuration file.
Press BACKUP to backup current configuration files.
Press RESTORE FROM SAVED to restore a previous configuration.
Press RESTORE FROM NEW to return files to unused condition.
Edit Configuration File
Backup...
Restore from saved...
Restore from new
Exit
Help

Figure 2-16. Accessing the Data File Manager Window

The Data File Manager window provides access to each of the Data File Manager functions: **Edit Configuration File**, which allows you to select the files and/or file types that will be backed up; **Backup**, which performs the backup action; **Restore from saved**, which replaces the current set of files with a previously-backed-up version; and **Restore from new**, which replaces all listed files with their original default versions.

The following sections detail how to perform each operation.

Editing the Backup Configuration File

When you perform a backup action, the Data File Manager looks in five directories in the \NetSight Element Manager x.x directory — \Data, \Bin, \Filters (and its automatically-created subdirectories), \Maps (or your configured Maps directory), and \Exports (or your configured Exported Files directory) — for the files and/or file types specified in the Backup Configuration File. Any files you wish to back up must be either specifically named (Tools.dat) in the Backup Configuration File or of a type (*.map) listed in the file; files must also be present in one of the five directories.

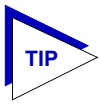

*If you change the default Maps and/or Exported Files directory (via the **Directories Options** available in the **Tools—>Options** window), the Data File Manager will automatically back up any specified files that reside in the new default directories. Any files that reside in the original default directories, however, will no longer be backed up.*

To view or edit the list of data files that will be backed up, click on **Edit Configuration File**. A text file display (stored as **\NetSight Element Manager x.x\bin\emanfil.db**), showing the current list of files and file types that will be backed up (Figure 2-17).

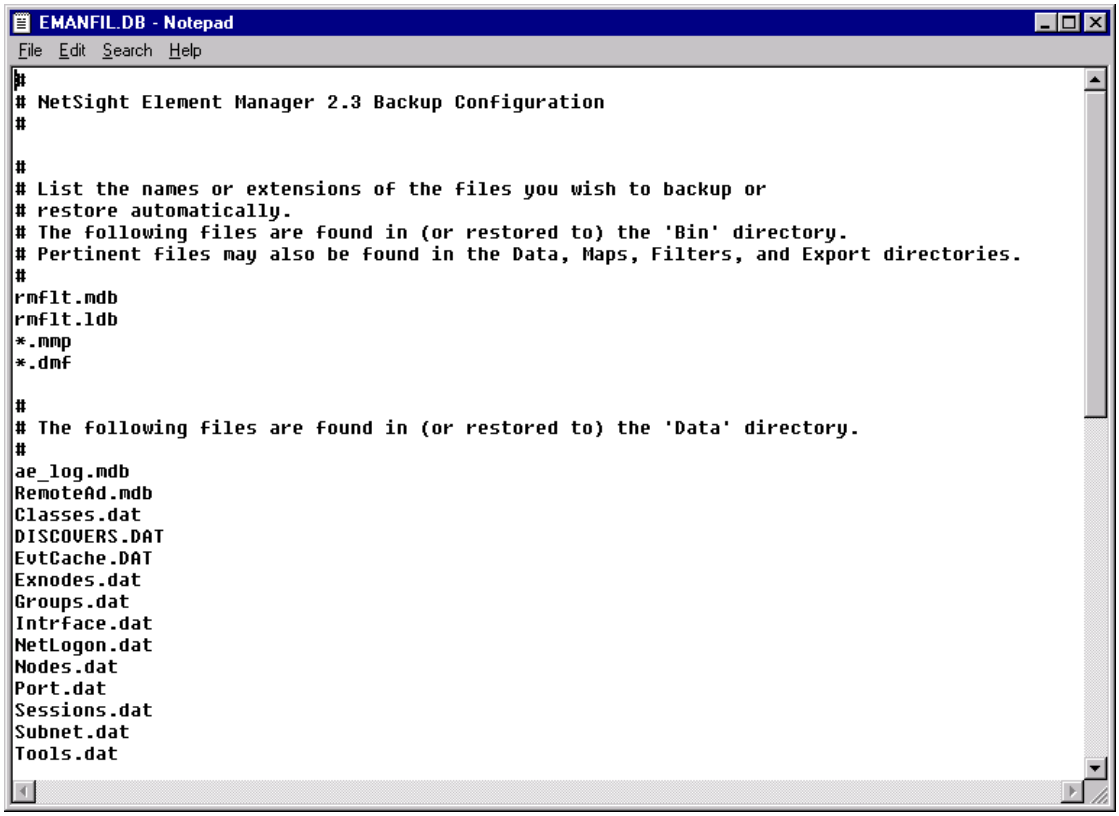

Figure 2-17. Sample NetSight Element Manager Backup Configuration

To edit this list:

1. Create a file backup entry by using the text editor to type a new line that specifies the file name, or the *.XXX file extension (for all files of a similar type). Remember, only files in the five pre-set directories will be backed up, so you needn't specify a path for any files you add to the list.
2. Save and exit the text file.

We encourage you to add the following files to your Backup Configuration file, if applicable:

- Any data files in your \Bin directory that have been placed as Linked Objects in a map (e.g., sound recorder files, Paintbrush files, database files, etc.)

Performing a Backup

1. In the main Data File Manager window, click on **Backup**. A window opens (Figure 2-18), prompting you to specify the directory path in which to copy the backup files.

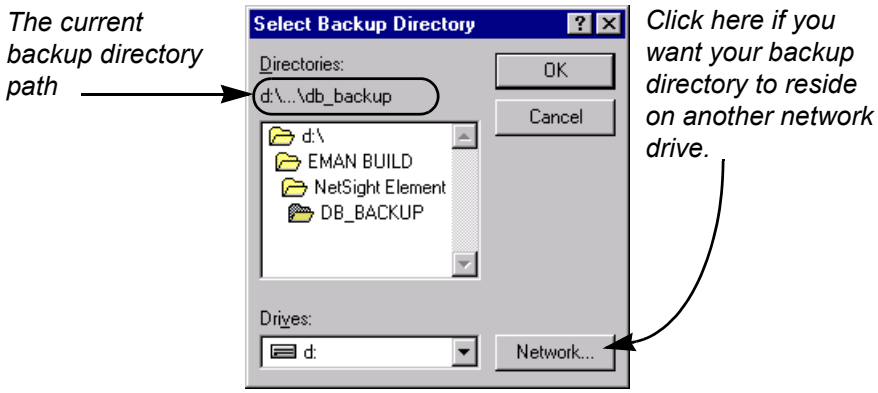

Figure 2-18. Select Backup Directory Window

NOTE

For best results, we recommend that you store a copy of your backed-up files in a directory that is not a sub-directory of your NetSight Element Manager installation directory.

2. Click **OK**. As the backup proceeds, you will be prompted to authorize the overwriting of any existing files. If any part of the procedure fails (for example, if you attempt to backup a specific file that does not reside in one of the five preset directories), you will be notified of the error.

3. When the backup is complete, a successful backup message will be displayed. Click **OK** to exit this window and complete the backup procedure.

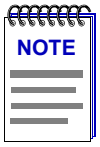

If even one file listed in the Backup Configuration file could not be found, a message window opens noting that the backup operation failed. This doesn't mean the entire operation failed; it just indicates the failure to back up the missing file.

Restoring Files from a Backup Set

If you have backed up your data files from SPECTRUM Element Manager version 2.x before installing NetSight Element Manager version 3.0, ***do not restore these data files directly into version 3.0****! The 3.0 installation process performs some necessary conversion functions on these files to make them compatible with framework changes. If you attempt to use version 2.x data files that have not undergone this conversion process will produce unexpected behavior. As soon as you have installed version 3.0 and verified that your files are working correctly, we recommend that you do another backup of your 3.0 data files.*

1. In the main Data File Manager window, click on **Restore from saved**. A window opens, prompting you to specify the directory path from which to copy the archived files back into their original locations. This will be the same directory path you chose in the steps described above.
2. Click **OK**. Two pop-ups open in quick succession: one to indicate the files are being backed up, and a second to indicate the procedure was successful. You will be given the opportunity to cancel the operation if you are overwriting an existing file or, if the procedure fails, you will be notified of the error.
3. Click **OK** to exit the final prompt window.

Restoring Default Files

To restore your configuration files to their original condition without re-installing:

1. Click on **Restore from new**. A message appears, reminding you that this operation will destroy some or all of your data files, and reset much of your configuration to default values.
2. Click **OK**. Two pop-ups open in quick succession: one to indicate that the files are being restored, and a second to indicate the procedure was successful. You will be given the opportunity to cancel the operation if you are overwriting an existing file or, if the procedure fails, you will be notified of the error.
3. Click **OK** to exit the final prompt window.

Chapter 3

Discovering Nodes

Accessing the Discover Manager window; creating and modifying Discover scripts; executing Discover scripts; the Discover process; starting the Discover Wizard; launching the Scheduler application; using the IP or Host Locator; using the Subnets window

---

The Discover Manager application allows you to locate devices on a network. You can perform a Discover process within a specified range of IP (Internet Protocol) addresses, on an entire subnet, or on a combination of the two. The Discover application automatically explores the defined network segment, creates a list of discovered devices, and assigns an icon to each of them. The discovered nodes are then automatically integrated into NetSight Element Manager's central node database, and consequently displayed as appropriate in any List or Tree Views (see Chapter 4, **List Views**, and Chapter 5, **Tree Views**, for more information).

The Discover application operates using scripts, which are created and saved within the application (allowing them to be run repeatedly if desired). You can create a new Discover script by using either the Discover Manager window (Figure 3-1) or the Discover Wizard, which will lead you step by step through the Discover process (see **Starting the Discover Wizard** on page 3-19). The following sections describe the use of the Discover Manager window to create, modify, and execute Discover scripts.

Accessing the Discover Manager Window

To access the Discover Manager window from the NetSight Element Manager primary window:

1. Select **Tools —> Discover Manager** from the primary window menu bar.

   or

   Click the Discover Manager button ( ) on the NetSight Element Manager Standard toolbar.

   The Discover Manager window, Figure 3-1, opens.

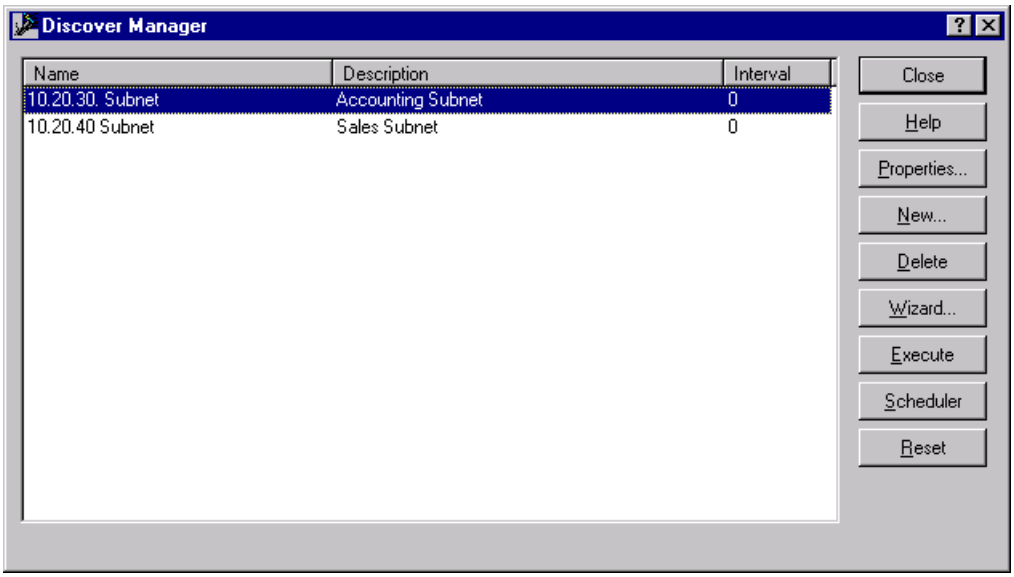

Figure 3-1. The Discover Manager Window

The Discover Manager window lists all Discover scripts you have created, including their names, descriptions, and execution intervals. The Discover Manager window is the central interface through which you create new scripts and modify, execute, and delete existing scripts; from this window, you can also choose to run the Discover Wizard, and launch the Scheduler application. (For more information about the Scheduler application, see the ***Tools Guide***.)

Creating and Modifying Discover Scripts

Discover scripts are created and modified using the Discover Properties window (Figure 3-2). When you are creating a new Discover script, this window will supply three tabbed pages — General, Device, and Schedule — which allow you to specify all desired parameters.

Accessing the Discover Properties Window

To create a new Discover script:

1. Click **New**.

   or

   Click the **right** mouse button on an existing script name or in the central area of the window, and choose **New** from the resulting menu.

The Discover Properties window, Figure 3-2, opens.

To modify an existing Discover script:

1. In the Discover Manager window, highlight the script to be modified, and click **Properties**.

   or

   Click the **right** mouse button on the script to be modified, and select **Properties** from the resulting menu.

The Discover Properties window, also illustrated in Figure 3-2, opens.

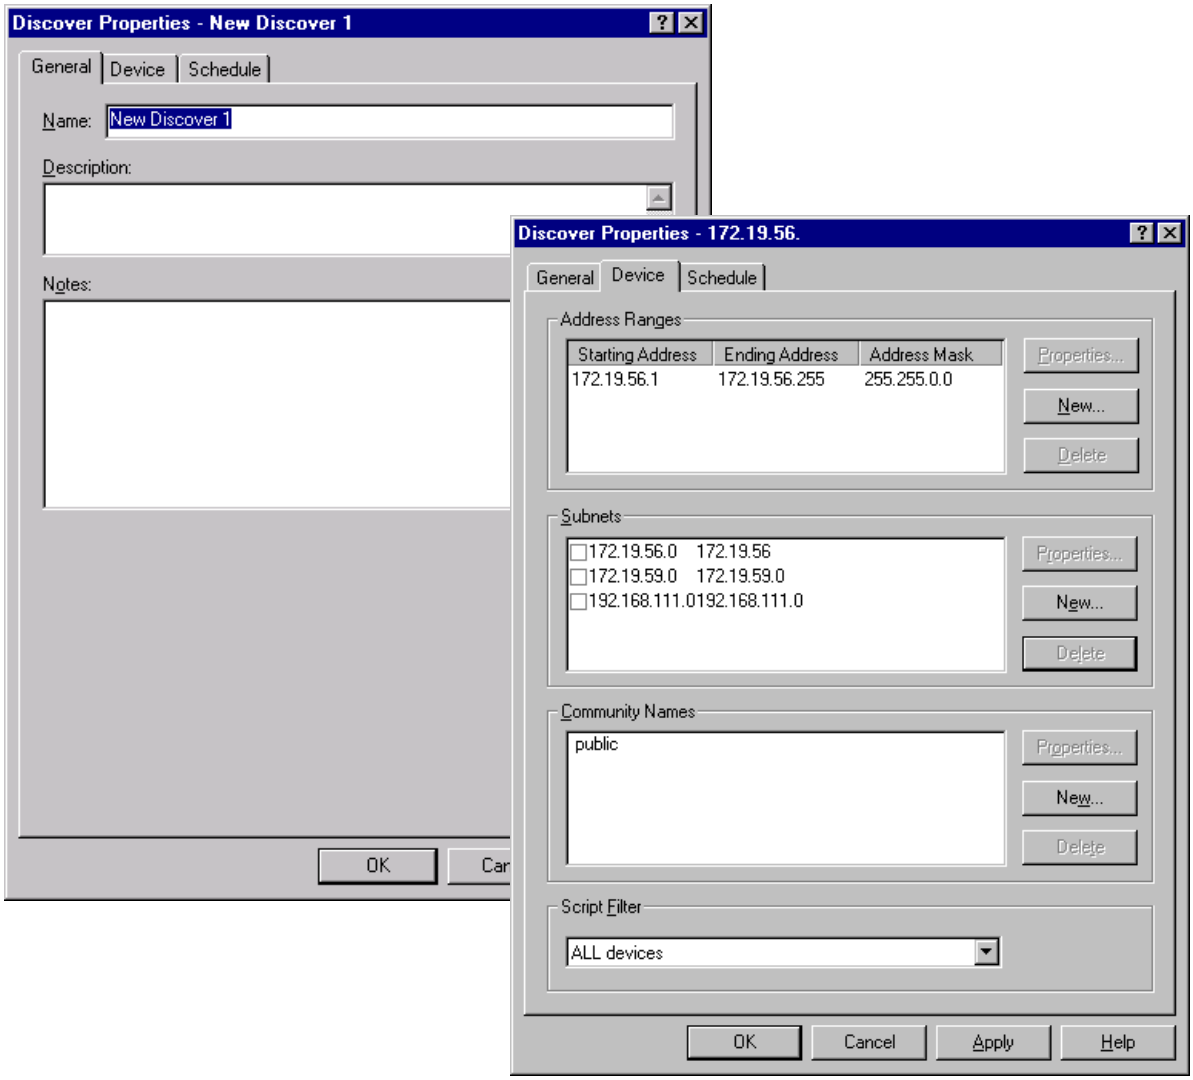

Figure 3-2. The Discover Properties Window

Using the Discover Properties window, you can name your script, enter an IP range for an IP Discovery process, and assign a community name. You can also define and select entire subnets for a Subnet Discovery, select a Script Filter for a discovery, and define the time interval at which your discovery will take place.

Setting General Discover Properties

The **General** page of the Discover Properties window is used to name and describe your Discover script (as shown in Figure 3-3). The script's name and description will be listed in the Discover Manager window.

To name and describe your Discover script:

1. If necessary, click on **General** tab in the Discover Properties window. The **General** page (Figure 3-3) opens.
2. Highlight the contents of the **Name:** field, and type the name of your Discover script. By default, each new script is named **New Discover X**, where X increments as necessary to create a unique name.

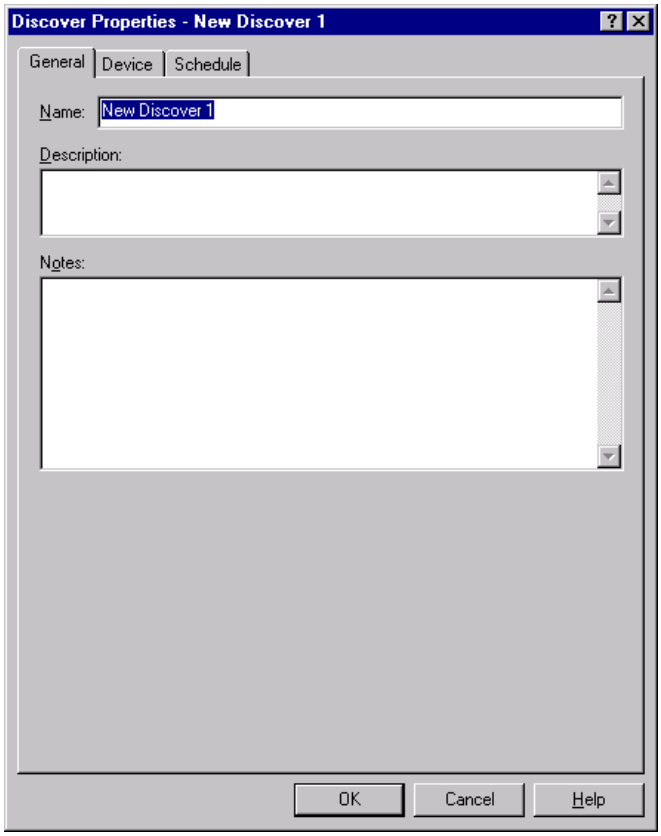

Figure 3-3. The **General** Page of the Discover Properties Window

3. In the **Description:** field, enter a description of the Discover script. This description opens with the script's name in the Discover Manager window, and may help you to select from among the available scripts.
4. Enter any additional information you wish to store about the script in the **Notes:** field. Note that this information is not displayed anywhere else.
5. Define your remaining Discover script parameters as explained in the following sections, or click **OK** to save your entries and return to the Discover Manager window.

Setting Device Discovery Parameters

A Device Discover script can discover devices within a single IP address range or a series of ranges. It can also search one or more subnets, and limit its discoveries to SNMP devices only or Cabletron devices only. You define Device Discovery parameters using the Device page of the Discover Properties window (Figure 3-4), accessed by clicking on the **Device** tab.

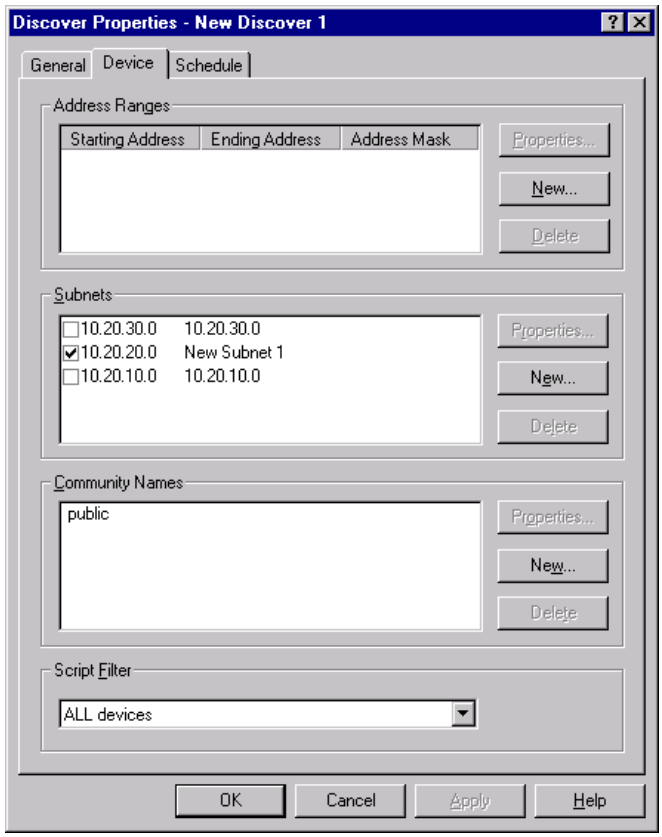

Figure 3-4. The **Device** Page of the Discover Properties Window

To define a Device Discovery script, you must specify a range of IP addresses to be queried and assign at least one community name that will give the Discover application read access to the devices you wish to discover. IP address ranges can be individually specified, or entire subnets can be selected; you can also combine subnets and specific ranges in a single script, as long as the total number of addresses to be queried does not exceed 2000. If you choose, you can also assign a Script Filter that will limit the discovery to SNMP devices only or Cabletron devices only.

Defining IP Address Ranges

To set the IP address range (or ranges) for your Device Discovery:

1. **To define a new IP address range**, click **New** to the right of the Address Ranges group box; or click the **right** mouse button anywhere in the Address Ranges group box, and select **New** from the resulting menu.

   To modify an existing IP address range, highlight the IP Address range that you wish to change, and click **Properties**; or click the **right** mouse button on the IP address range that you wish to modify, and select **Properties** from the resulting menu.

   The Address Range Properties window (Figure 3-5) opens.

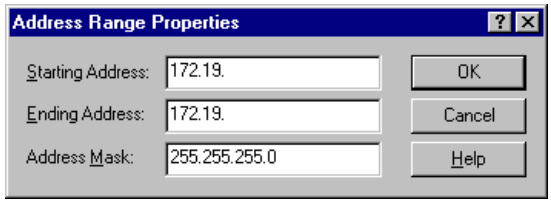

Figure 3-5. The Address Range Properties Window

NOTE

*When you are defining a new address range, the **Starting Address**, **Ending Address**, and **Address Mask** fields will contain the autofill values you configured via the Welcome Wizard or via the **Network** page in the **Tools—>Options** window (provided the autofill function has been enabled via the **Workspace** page of the **Tools—>Options** window). See **Network Options** in Chapter 2 for more information about configuring these defaults.*

2. Place the cursor in the **Starting Address** field, and enter the address at which your Discover script should begin.

3. Place the cursor in the **Ending Address** field (using the **Tab** key or the mouse), and enter the address at which the Discovery should end.

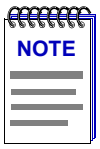

Although you can define address ranges beginning with 0 and ending with 255 (for example, 132.117.18.0 to 132.117.18.255), devices at the .0 and .255 addresses will not be discovered.

4. If you wish to specify an address mask other than the default value, enter the desired address mask in the **Address Mask** field.

   The appropriate address mask depends upon your network configuration. If your network contains subnets, entering the appropriate *subnet mask* in this field will identify the network, subnet, and host portions of the IP addresses in the address range. For example, on a Class B network containing an IP address of 132.177.118.24, a subnet mask of 255.255.255.0 would identify the first two octets (132.177) as the network portion of the address, the third octet (118) as the subnet portion, and the final octet (24) as the host portion.

   If your network does not contain subnets, entering the appropriate *network mask* in this field will identify the network portion of the IP addresses in the address range. For example, on a Class B network containing an IP address of 132.177.118.24, a network mask of 255.255.0.0 would identify the first two octets (132.177) as the network portion of the address.

5. Click **OK** when your entries are complete. The information you entered displays in the appropriate fields in the Address Ranges group box.

6. Repeat this process to define additional address ranges for your script.

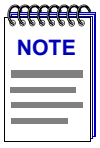

Each discover script can be configured to query up to 2000 IP addresses; if the number of IP addresses on your network exceeds 2000, you must define and execute multiple Discover scripts.

7. Define your remaining Discover script parameters as explained in the following sections, or click **OK** to save your entries and return to the Discover Manager window.

Defining and Selecting Subnets

Using the Subnets group box, you can also select entire subnets on which to perform a Discovery. When a Subnet Discovery is executed, the Discover application will query every address in the specified subnet, making it unnecessary to define an IP Address range that encompasses an entire subnet.

Subnets will be defined automatically each time a device is added to the node database (either manually or via a Discover script); they can also be defined manually, either via the Discover Properties window (as described in this section) or via the **View—>Subnets** window available from the NetSight Element Manager primary window menu bar. (See **Using the Subnets Window**, page 3-27, for more information about viewing subnets.)

This subnet information is used to determine the map hierarchy created via the automatic subnet mapping feature; for more information on this feature, see ***Generating Submaps Automatically*** *in Chapter 6,* ***Creating Network Maps****.*

To specify a subnet (or subnets) for your device discover script:

1. **To define a new subnet**, click **New** to the right of the Subnets group box; or click the **right** mouse button anywhere in the Subnets group box, and select **New** from the resulting menu.

   To modify an existing subnet, highlight the subnet entry that you wish to change, and click **Properties**; or click the **right** mouse button on the subnet entry that you wish to modify, and select **Properties** from the resulting menu.

   The Subnet Properties window (Figure 3-6) opens.

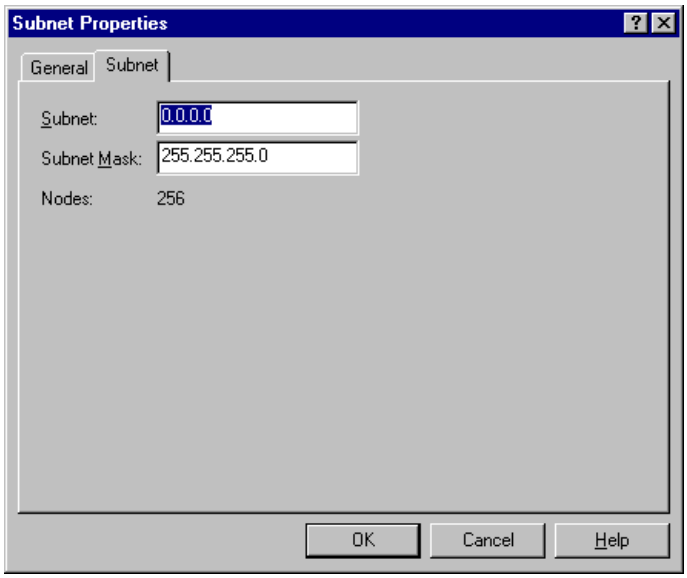

Figure 3-6. The **Subnet** Page of the Subnet Properties Window

2. Place the cursor in the **Subnet** field and type the network and subnet portions of the IP Address for the subnet that you wish to discover. For example, to discover the subnet containing the Class B IP Address 132.177.118.24, type 132.177.118.0 in the **Subnet** field).

3. If necessary, place the cursor in the **Subnet Mask** field and type the subnet mask for the subnet that you wish to discover. Note that the value set here will determine the number of IP addresses which are included in the defined subnet, regardless of whether or not this value correctly reflects the number of devices each of your subnets contains. For example, a Class B network typically defines the third byte of the IP address as the subnet identifier, and the fourth byte as the host identifier. In this kind of network, each subnet (as defined by each possible value of the third byte) can contain a maximum of 256 nodes on it (the number of possible values for the fourth byte). However, if you incorrectly identify the subnet mask for such a network as 255.255.0.0, the discover application interprets this to mean that each subnet has 65,536 IP addresses (256 possible values of the third byte, times 256 possible values for the fourth byte). Since each discover script is limited to querying a maximum of 2000 addresses, this kind of incorrect mask setting would prevent you from completing your script. By default, the Subnet Mask will be set to 255.255.255.0, limiting the size of the subnet to a maximum of 256 devices. The **Nodes** field just below the **Subnet Mask** field indicates the number of nodes the Discover application understands your defined subnet to contain; it will update once you have saved your subnet changes, closed the Subnet Properties window, and re-opened it.

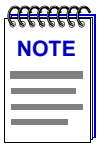

*The **Nodes** field in the Subnet Properties window is not the same as the nodes value displayed in the **View—>Subnets** window (see Figure 3-20, page 3-28); that value indicates the actual number of nodes in the central node database that reside on the defined subnet.*

4. On the **General** page of the Subnet Properties window (Figure 3-7), place the cursor in the **Name** text box, and type in a name for your subnet entry. By default, manually-created subnets are assigned the name **New Subnet X**, where X will increment as necessary to create a unique name. Automatically-created subnets are assigned the subnet address as their default name.

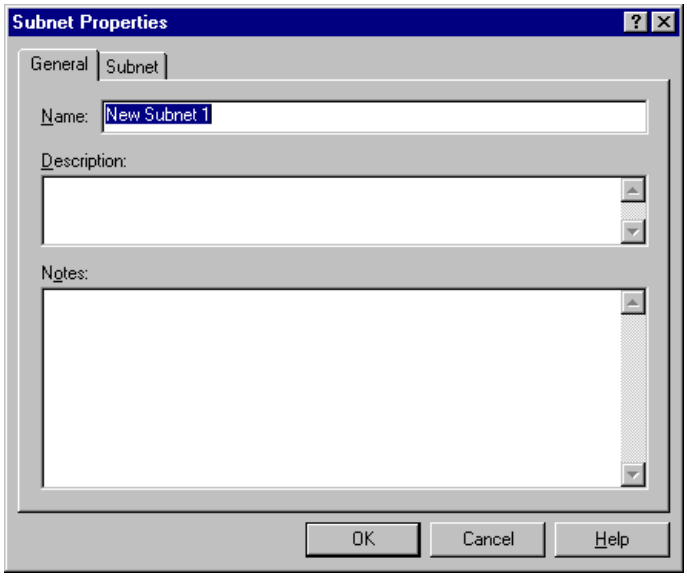

Figure 3-7. The **General** Page of the Subnet Properties Window

5. In the **Description** text box, enter a description for your subnet entry. This description is not displayed in the Discover Properties window, but will be displayed in the **View—>Subnets** window accessible from the primary window menu bar (see Figure 3-20 on page 3-28).
6. Enter any additional information about the subnet in the **Notes** text box.
7. Click **OK** to save your subnet discovery parameters and exit the Subnet Properties window. The address and name of the defined subnet displays in the Subnets group box on the **Device** page of the Discover Properties window.
8. Repeat Steps 2-8 to define additional Subnets for your script.
9. To include a defined subnet in your Discover script, click on the checkbox to the left of the subnet address in the Subnets group box. A check mark displays in the checkbox to indicate that the subnet discovery has been selected for inclusion in the script.
10. Define your remaining Discover script parameters as explained in the following sections, or click **OK** to save your entries and return to the Discover Manager window.

TIP

You can combine one or more IP address ranges with one or more subnets in a single device discover script, as long as the total number of addresses to be queried in a single script does not exceed 2000.

Defining a Community Name

Once you have specified the IP addresses that will be queried during the Discover, you must supply the discover application with one or more community names that will provide at least Read access to the devices you wish to discover.

When a script with more than one assigned community name is executed, the Discover application will query each address in the specified address range(s) and/or subnet(s) using the first community name listed in the Community Names group box. If a node responds, the other community names are not used to query that node. If there is no response to the first query, the address is queried again using the second listed community name. This process is repeated until a response is returned or until all listed community names have been used to query the address. If no response is elicited from an address after using all community names assigned to the script, the Discover application gives up on that address, and moves on to the next one.

To select a community name for use in your Discover script:

1. **To define a new community name**, click **New** to the right of the Community Names group box; or click the **right** mouse button anywhere in the Community Names group box, and select **New** from the resulting menu.

   To modify an existing community name, highlight the community name that you wish to change, and click **Properties**; or click the **right** mouse button on the community name that you wish to modify, and select **Properties** from the resulting menu.

   The Community Name Properties window (Figure 3-8) opens.

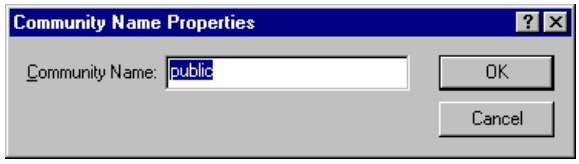

Figure 3-8. The Community Name Properties Window

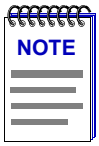

*If a default community name has been specified (either via the Welcome Wizard or via the **Network** page of the **Tools—>Options** window) and the autofill function has been activated (via the **Workspace** page of the **Tools—>Options** window), the Community Name properties window will contain a default community name value.*

2. If the **Community Name:** field contains a default name, it will be highlighted. If you wish to use a different community name, type the new name into the field.
3. Click **OK** to save the community name and close the window. The community name you have entered will be listed in the Community Names group box.

4. To add another community name to your script, repeat steps 2-4.
5. Define your remaining Discover script parameters as explained in the following sections, or click **OK** to save your entries and return to the Discover Manager window.

Selecting a Script Filter

If you would like to further refine a device discovery script, you can add one of two Script Filters: one which restricts the discovery to Cabletron devices, and one which restricts it to devices which respond to an SNMP query.

1. Click on the Script Filters drop-down menu; the three available filters will be displayed, as illustrated in Figure 3-9.

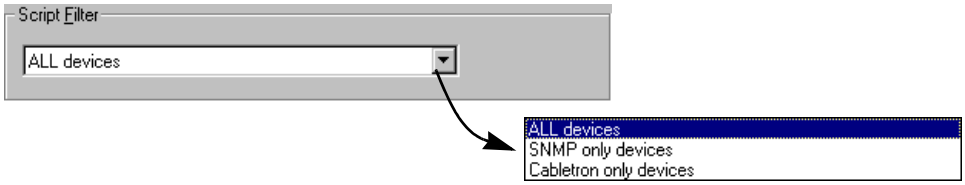

Figure 3-9. Script Filter Options

2. Click to select the appropriate filter:

   All Devices — the script will discover and add to the central node database all devices in the specified subnet(s) or IP address range(s) that respond to either an SNMP or a PING query. This is the default value.

   SNMP only devices — the script will discover and add to the central node database only those devices in the specified subnet(s) or IP address range(s) that respond to an SNMP query. Selecting this filter eliminates the discovery of “Ping” devices.

   Cabletron only devices — the script will discover and add to the central node database only those devices in the specified subnet(s) or IP address range(s) that respond to an SNMP query and are identified as Cabletron devices. Selecting this filter eliminates the discovery of third-party devices.

3. Define your remaining Discover script parameters as explained in the previous sections, or click **OK** to save your selection and return to the Discover Manager window.

Setting a Discover Interval

Defining a Discover Interval for your Discover script allows you to automatically start your script and determine how often (in minutes) your script will be automatically executed. The Discover Interval is defined using the **Schedule** page of the Discover Properties window (Figure 3-10), accessed by clicking on the **Schedule** tab.

Once you have defined an interval and saved the script, NetSight Element Manager will automatically execute the script, and continually re-execute it according to the defined time interval. The script will continue to run at the defined interval until its Discover Interval is changed to 0 (zero), the script is deleted, or the NetSight Element Manager application is closed.

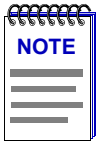

If you define a script with a Discover Interval, and then exit NetSight Element Manager, the script will be automatically executed the next time you start NetSight Element Manager.

To set an interval for your Discover script:

1. In the **Discover Interval** field, enter the desired number of minutes between automatic executions of your Discover script. This interval will define the amount of time between the *beginning* of one discover session and the *beginning* of the next one (so be sure you define an interval that allows enough time for each session to complete). The default value of 0 (zero) indicates that manual execution of the script is required.
2. Define any remaining Discover script parameters, or click **OK** to save your interval and return to the Discover Manager window.

Note that scheduled discovers behave a little differently than manually-executed discovers; see **The Discover Process**, page 3-16, for more information.

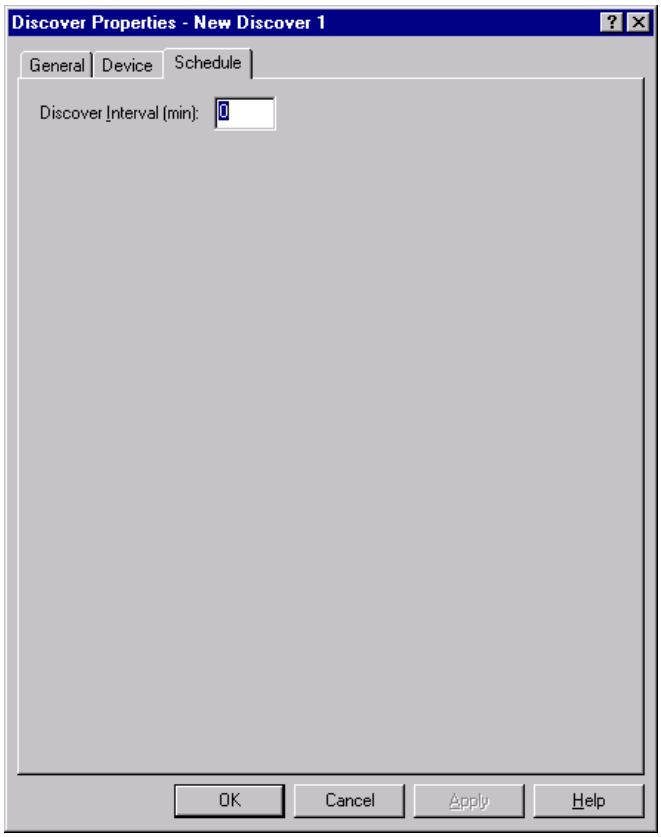

Figure 3-10. The **Schedule** Page of the Discover Properties Window

Executing Discover Scripts

Once you have established the parameters for your Discover scripts, they are listed in the Discover Manager window (as shown in Figure 3-11), and can be executed at any time.

1. In the Discover Manager window, click to highlight the name of the Discover script to be executed, and click **Execute**.

   or

   In the Discover Manager window, click the **right** mouse button on the name of the desired Discover script to reveal the Discover Manager menu, and select **Execute**.

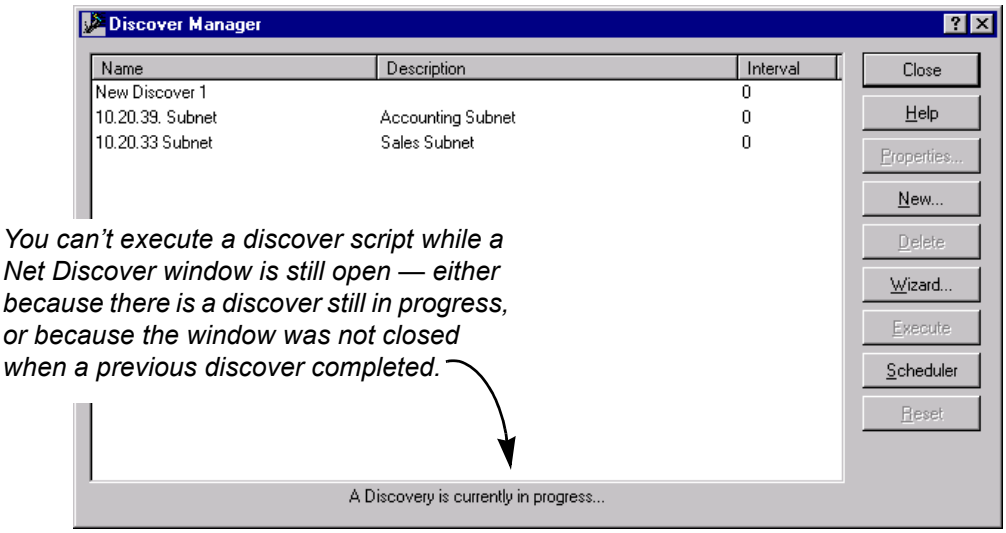

Figure 3-11. The Discover Manager Window

You cannot execute a new discover session if a previous session is still running or if the Discover window from a previous session is still open. If either of these conditions exists, a footer message at the bottom of the Discover Manager window will indicate that a discover is in progress (as illustrated in Figure 3-11). To clear this condition, close the Discover Manager window, activate the Network Discover window (which may be minimized or hidden behind the main NetSight Element Manager window), wait for the active discover session to complete (or stop it, as described in **Stopping a Discover in Progress**, page 3-17), close the Network Discover window, re-launch the Discover Manager, and execute your new session.

When you have successfully launched a discover session, the Discover Manager window automatically closes, and the Network Discover window (Figure 3-12, below) opens. For more detailed information about this window and the discover process, see the following sections.

The Discover Process

Once a Discover script has been manually executed, the Discover Manager window closes, and the Network Discover window (Figure 3-12) opens.

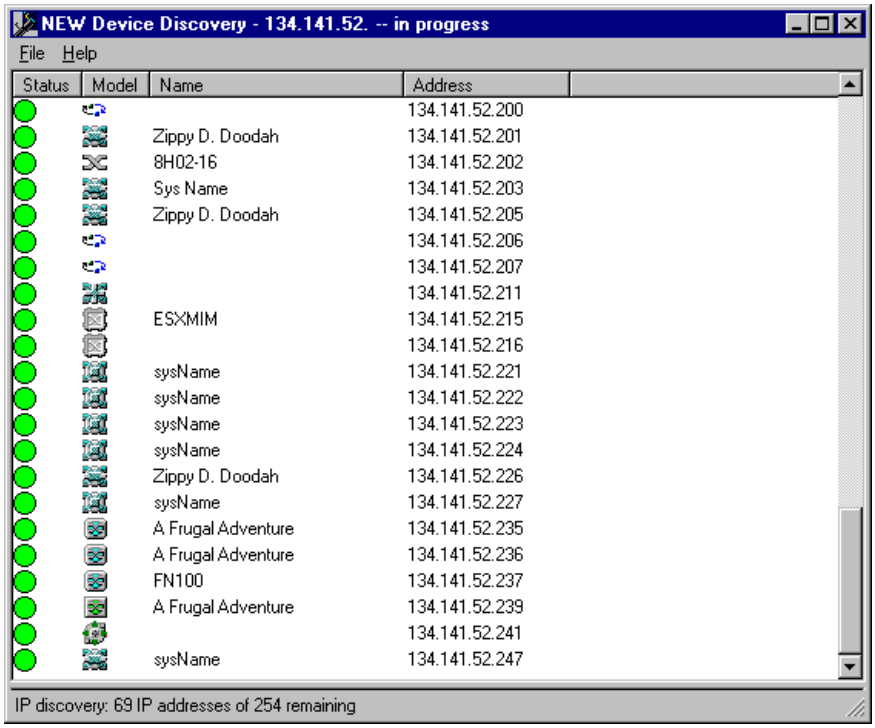

Figure 3-12. Network Discover Window

A discover proceeds in a single phase: each IP address in the defined address ranges and/or subnets is queried in turn; as each response is received, information about the responding device (including status, model, name, and IP address) appears in the Discover window. A footer message at the bottom of the window indicates the total number of IP addresses which will be queried during the discover session, and the number of IP addresses which remain to be contacted.

When the Discover process is complete, all discovered devices which match the configured criteria are automatically added to NetSight Element Manager's central node database. A progress indicator shows the status of this automatic integration; the Discover window also remains open after the discover is complete, allowing you to view the list of discovered nodes.

NEW and SECONDARY Discovers

Discovery sessions are divided into two types: NEW and SECONDARY. The first time you run a device discover script, all IP addresses in the subnet or address range you have specified will be queried, and all discovered devices will be listed in the discover window and automatically added to the central node database. This is called a NEW Discovery. The *second* time you run any device discover, however — even a device script which has not yet been run — NetSight Element Manager will perform a SECONDARY Discover, querying only those IP addresses which did not respond to the first discover.

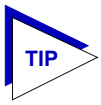

When a NEW Discovery is performed, new information which may be discovered about an IP address will not overwrite the information about that address which is currently in the central node database. For example, if an IP address previously assigned to an EMME is moved to a SmartSwitch 6000, the discover application does not record this change in the database.

Stopping a Discover in Progress

1. Select **File—>Exit** from the Network Discover window's menu bar.

   or

   Click the close button in the upper-right corner of the Network Discover window. A message window like the one illustrated in Figure 3-13 opens.

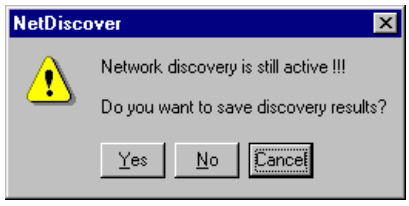

Figure 3-13. Cancel Discover Confirmation Window

2. Click **Yes** to stop the discover and add any discovered nodes to the database; click **No** to stop the discover without adding any partial results to the database; or click on **Cancel** to continue the discover.

Restarting a Completed Discover

You can also re-start a completed discover from the Discover window; note that this will launch the selected script in SECONDARY discover mode, unless the discover counter has been reset.

1. In the Network Discover window, select **File—>Rediscover**. The original discover script will be re-launched, in SECONDARY discover mode.

Re-starting a completed discover script from the Network Discover window has the same effect as re-executing the script from the Discover Manager window.

Scheduled Discovers

A scheduled discover script — one which launches automatically at selected intervals — behaves a little differently than a manually-executed script. Once a scheduled discover is complete, the Network Discover window will close automatically; you need not close it manually, as with a normally-executed script. (In fact, the Network Discover window for a scheduled discover is minimized by default, so it is only visible if you choose to view it.) If you have left a manually-executed discover window open, a scheduled discover will close this window for you so that it can execute. And if a manually-executed script is still in progress when a scheduled discover interval expires, the scheduled discover will wait the length of the interval before it attempts to launch again.

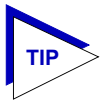

*You can also schedule discover sessions at specific dates and/or times via the Scheduler application; see the **Tools Guide** for more information.*

Deleting a Discover Script

1. In the Discover Manager window, highlight the name of the Discover script to be deleted, and click **Delete**.

   or

   In the Discover Manager window, **right-click** on the name of the Discover script to be deleted, and select **Delete** from the resulting menu.

   A delete confirmation window (Figure 3-14) opens.

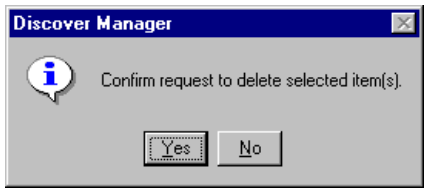

Figure 3-14. The Script Deletion Confirmation Window

2. Click **Yes** to confirm the deletion. The script will be removed from the Discover Manager window.

Starting the Discover Wizard

To use the Discover Wizard *to create a new Discover script*:

1. In the Discover Manager window, click **Wizard**.

   or

   Click the **right** mouse button anywhere in the Discover Manager window, select **Wizard** from the resulting menu.

To use the Discover Wizard *to modify an existing Discover script*:

1. In the Discover Manager window, highlight the name of the script that you wish to modify, and click **Wizard**.

   or

   Click the **right** mouse button on the name of the script that you wish to modify, and choose **Wizard** from the resulting menu.

Follow the directions in the Wizard to create or modify a discover script.

Launching the Scheduler Application

Once you have created your Discover scripts, you can use the Scheduler application to automatically execute selected scripts at scheduled times and dates. To launch the Scheduler application from the Discover Manager window:

1. Click **Scheduler**.

   or

   Click the **right** mouse button anywhere in the Discover Manager window, and choose **Scheduler** from the resulting menu.

For more information about using the Scheduler application, consult the ***Tools Guide***, included with your documentation.

Using the IP or Host Locator Tool

The IP or Host Locator Tool serves as an adjunct to the Discover functionality. Its primary function is to resolve host names to IP addresses, and vice versa; it can also resolve host names to IP addresses for an entire subnet, locate the MAC address associated with a matched host name/IP pair, and locate both the host name and IP address for a given MAC address.

To launch the IP or Host Locator Tool:

1. Select **Tools —>Host/IP Locator** from the primary window menu bar.

or

Click the **IP or Host Locator** button ( ) on the NetSight Element Manager Standard toolbar.

The IP or Host Locator window, Figure 3-15, opens.

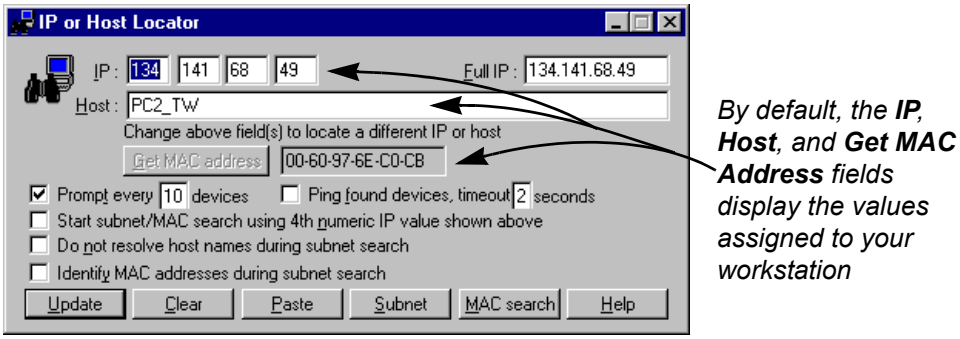
IP or Host Locator
IP : 134 141 68 49
Full IP : 134.141.68.49
Host : PC2_TW
Change above field(s) to locate a different IP or host
Get MAC address
00-60-97-6E-C0-CB
Prompt every 10 devices
Ping found devices, timeout 2 seconds
Start subnet/MAC search using 4th numeric IP value shown above
Do not resolve host names during subnet search
Identify MAC addresses during subnet search
Update
Clear
Paste
Subnet
MAC search
Help
By default, the IP, Host, and Get MAC Address fields display the values assigned to your workstation

Figure 3-15. The IP or Host Locator Window

The following sections describe some common uses for this tool. Note that multiple instances of the IP or Host Locator tool can be run simultaneously, allowing you to perform more than one lookup task at a time.

Matching a Host Name to an IP Address

To use the Locator Tool to locate the IP address associated with a known host name:

1. Click **Clear** to clear the default values from the IP and Host fields.
2. In the **Host** field, enter the host name of the device whose IP address you wish to locate.
3. If you wish to ping the device once its IP address is located, click to select the **Ping found devices** option; in the text field provided, specify a timeout period, in seconds. If the device does not respond to the ping request within the specified timeout, the ping will be considered to be unsuccessful.
4. Click on **Update**, or press Enter on the keyboard. If the host name can be resolved, its associated IP address will be displayed in the IP address fields; if you have selected the **Ping** option, a message window will indicate the success or failure of the ping attempt. If the host name cannot be resolved, a message window appears, indicating that the selected host name could not be found.

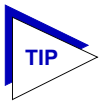

*Once located, IP addresses can be copied (using the right-mouse edit menu or **Ctrl-C** on the keyboard) and pasted (via **Edit—>Paste**) into the Insert Device window for addition to the central node database. For more information about inserting devices individually, see Chapter 4,* **List Views**.

Matching an IP Address to a Host Name

You can also use a known IP address to locate a device's host name. To do so:

1. Click **Clear** to clear the default values from the IP and Host fields.
2. Enter the IP address of the device whose host name you wish to locate in the **IP** or **Full IP** field.
3. If you wish to ping the device once its host name is located, click to select the **Ping found devices** option; in the text field provided, specify a timeout period, in seconds. If the device does not respond to the ping request within the specified timeout, the ping will be considered to be unsuccessful.
4. Click on **Update**, or press Enter on the keyboard. If the IP address can be resolved, its associated host name will be displayed in the **Host** field; if you have selected the **Ping** option, a message window will indicate the success or failure of the ping attempt. If the IP address cannot be resolved, a message window appears, indicating that the selected IP could not be found.

Locating the MAC Address for a Host Name/IP Pair

Once you have resolved a host name/IP pair, you can use the **Get MAC address** button to locate the device's MAC address.

1. Enter an IP or host name in the appropriate field, and click **Update** to resolve the pair.
2. Click on **Get MAC address**. The MAC address associated with the IP/Host pair will be displayed in the text box to the right of the button.

You cannot use an IP or host name alone to retrieve a MAC address; you must first resolve the pair to activate the **Get MAC address** button.

If you have a MAC address and wish to find the associated IP and/or host name, use the MAC Search feature described in ***Matching a MAC Address to a Host Name and IP****, beginning on page 3-25.*

Subnet Search

You can also use the Locator Tool to locate (and ping, if desired) all IP addresses on a subnet or a portion of a subnet, resolve host names to the located IP addresses, and resolve MAC addresses to the located IP/hostname pair. To do so:

1. Click **Clear** to clear the default values from the IP and Host fields.
2. In the **IP** field, enter the appropriate bytes of the subnet address:

   To search an entire subnet, enter the first three bytes of the subnet address;

   To search a portion of a subnet, enter the first three bytes of the subnet address; in the fourth byte field, enter the IP at which you want the search to start.
3. If you have chosen to search only a portion of a subnet, click to select the **Start subnet/MAC search using 4th numeric IP value shown above** option. A check mark opens in the selection box.
4. Because searching an entire subnet can be time-consuming, you can configure the Locator Tool to stop at pre-defined intervals, asking if you want to continue the search. To activate this option, click to select **Prompt every XX devices**, and use the text field provided to enter the number of devices that will be located between prompts. The default value is 10.

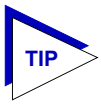

*Once the search begins, a **Stop Searching** button will also appear, allowing you to stop the search at any time.*

5. If you wish to ping each found device once it is located, select the **Ping found devices** option; in the text field provided, specify a timeout period, in seconds. If a device does not respond to the ping request within the specified timeout, the ping will be considered to be unsuccessful.
6. For the quickest Subnet search, click to select the **Do not resolve host names during subnet search** option. This kind of search will locate only IP addresses (and ping them if you have selected that option); no host name or MAC resolution will take place. Note that selecting this option inactivates the **Identify MAC address** option described in the next step.
7. For the most detailed Subnet search, click to select the **Identify MAC addresses during search** option. Though time-consuming, this subnet search gives you complete IP, host, and MAC resolution for all devices in the subnet (or subnet portion) you have specified.
8. Click **Subnet** to start the search. The Locator window will expand (see Figure 3-16) to display a list box containing the results of the search. A **Stop Searching** button will also appear just above the list box; click this button at

any time to stop a search in progress. If you have selected the **Prompt** option, a message window opens at the selected interval, asking if you want to continue the search. Click **Yes** to continue, or **No** to stop.

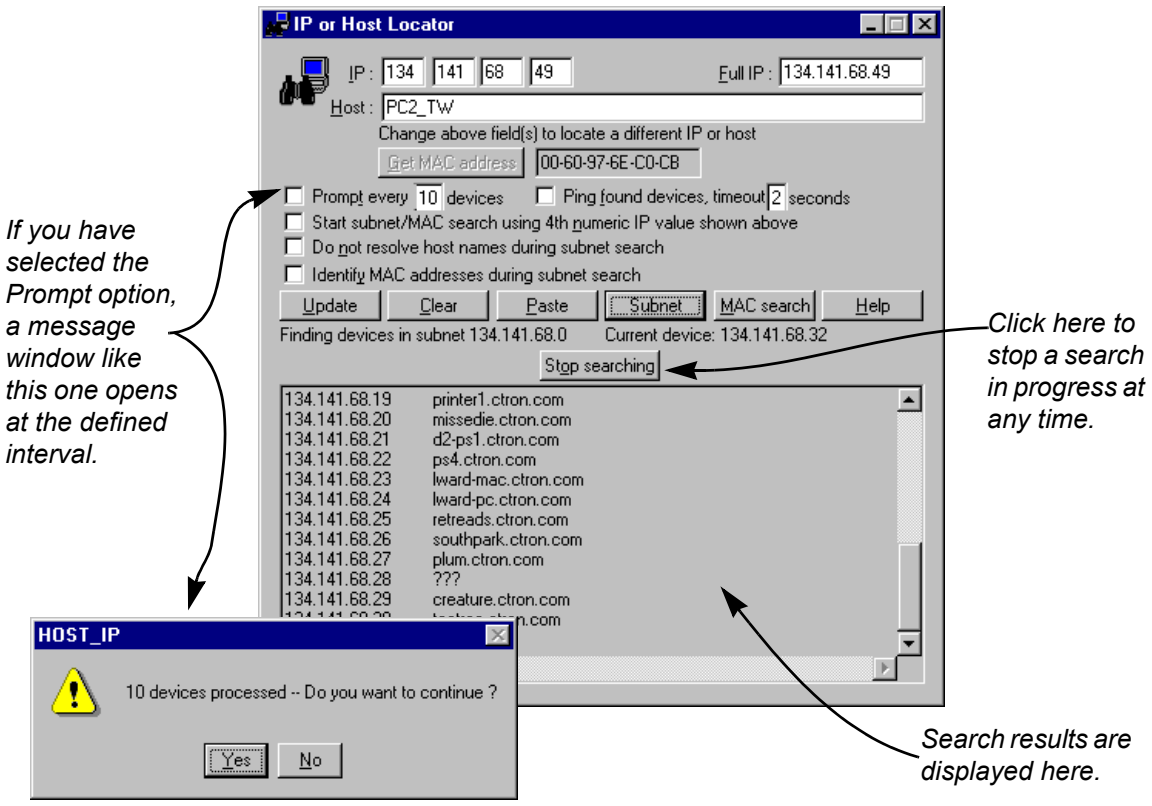

Figure 3-16. A Subnet Search in Progress

If you have selected the **Identify MAC address** and/or **Ping** options, these results will also be displayed in the list box. If no device is found at a specified IP address, question marks (???) display in place of a host name.

Once the search is complete, five buttons display just above the list box display (see Figure 3-17): **Select All**, **Copy IPs**, **Copy Unnamed IPs**, **Copy Pinged IPs**, and **Copy Unpingable IPs**.

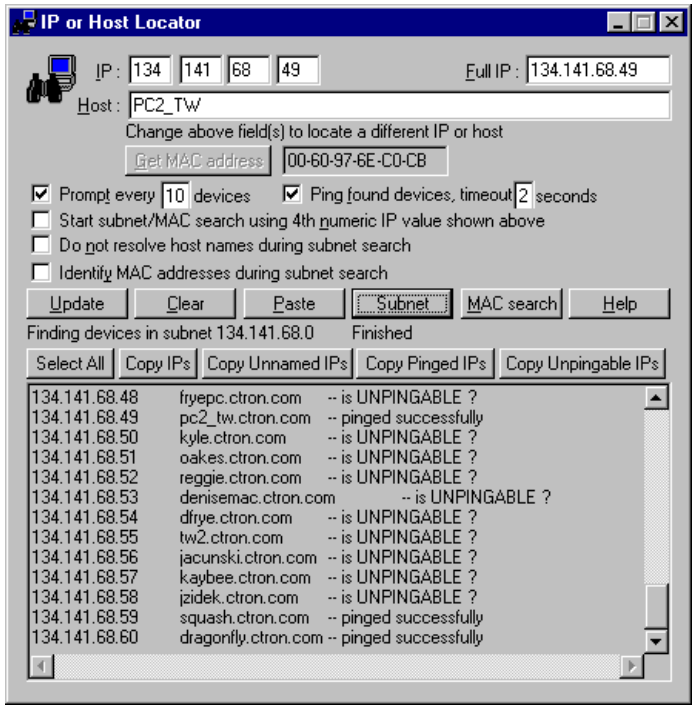
IP or Host Locator
IP : 134 141 68 49
Full IP : 134.141.68.49
Host : PC2_TW
Change above field(s) to locate a different IP or host
Get MAC address
00-60-97-6E-C0-CB
Prompt every 10 devices
Ping found devices, timeout 2 seconds
Start subnet/MAC search using 4th numeric IP value shown above
Do not resolve host names during subnet search
Identify MAC addresses during subnet search
Update
Clear
Paste
Subnet
MAC search
Help
Finding devices in subnet 134.141.68.0
Finished
Select All
Copy IPs
Copy Unnamed IPs
Copy Pinged IPs
Copy Unpingable IPs
134.141.68.48 fryepc.ctron.com -- is UNPINGABLE ?
134.141.68.49 pc2_tw.ctron.com -- pinged successfully
134.141.68.50 kyle.ctron.com -- is UNPINGABLE ?
134.141.68.51 oakes.ctron.com -- is UNPINGABLE ?
134.141.68.52 reggie.ctron.com -- is UNPINGABLE ?
134.141.68.53 denisemac.ctron.com -- is UNPINGABLE ?
134.141.68.54 dfrye.ctron.com -- is UNPINGABLE ?
134.141.68.55 tw2.ctron.com -- is UNPINGABLE ?
134.141.68.56 jacunski.ctron.com -- is UNPINGABLE ?
134.141.68.57 kaybee.ctron.com -- is UNPINGABLE ?
134.141.68.58 jzidek.ctron.com -- is UNPINGABLE ?
134.141.68.59 squash.ctron.com -- pinged successfully
134.141.68.60 dragonfly.ctron.com -- pinged successfully

Figure 3-17. Completed Subnet Search

You can use these buttons to copy one, several, or all discovered IP addresses to the clipboard for pasting into the central node database or any text document:

1. Click **Select All** to select all discovered IP addresses, or click and drag with your mouse to select a contiguous portion of the discovered address list.

   or

   Click **Copy Unnamed IPs**, **Copy Pinged IPs**, or **Copy Unpingable IPs** to select specific groups of IPs. (Note that these buttons only activate when IPs of the noted kind are displayed in the list box.)

2. Click **Copy IPs** to copy the selected IP addresses to the clipboard. Use any **Paste** command as usual to copy the contents of the clipboard into a text or other document, or use **Edit—>Paste** from the primary window menu bar to sequentially paste the copied IP addresses into a series of Insert Device windows for addition to the central node database.

NOTE

*Only the located IP address is copied during the **Copy IPs** process; no additional information, though highlighted in the list box display, is copied to the clipboard.*

Matching a MAC Address to a Host Name and IP

You can also use a known MAC address and subnet location to locate the associated IP address and host name. To do so:

1. Click **Clear** to clear the default values from the IP and Host fields.
2. In the **IP** field, enter the appropriate bytes of the subnet address:

   To search an entire subnet, enter the first three bytes of the subnet address;

   To search a portion of a subnet, enter the first three bytes of the subnet address; in the fourth byte field, enter the IP at which you want the search to start.
3. If you have chosen to search only a portion of a subnet, click to select the **Start subnet/MAC search using 4th numeric IP value shown above** option. A check mark displays in the selection box.
4. Because searching an entire subnet can be time-consuming, you can configure the Locator Tool to stop at pre-defined intervals, asking if you want to continue the search. To activate this option, click to select **Prompt every XX devices**, and use the text field provided to enter the number of devices that will be located between prompts. The default value is 10.

TIP

*Once the search begins, a **Stop Searching** button will also appear, allowing you to stop the search at any time.*

5. Click **MAC search**, and enter the MAC address whose IP and host name you wish to find in the window which appears (see Figure 3-18).

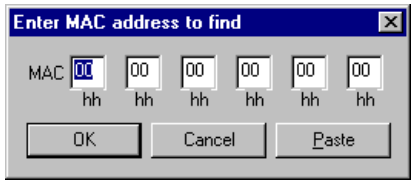

Figure 3-18. MAC Search Window

6. Click **OK** to start the search. The Locator window will expand (see Figure 3-19) and a **Stop Searching** button displays just above the list box; click this button at any time to stop a search in progress. If you have selected the **Prompt** option, a message window opens at the selected interval, asking if you want to continue the search. Click **Yes** to continue, or **No** to stop.

When the selected MAC address is located, a message window will display, indicating the host name and IP address associated with the found MAC; this information will also be displayed in the list box. If the address is not located, a message window opens when the search is complete, indicating that the selected address could not be found.

TIP

*You can use the **Copy IP** buttons to paste any located IP addresses to any text document; you can also paste found IPs directly into the central node database via the **Edit—>Paste** command. See page 3-24 for more information.*

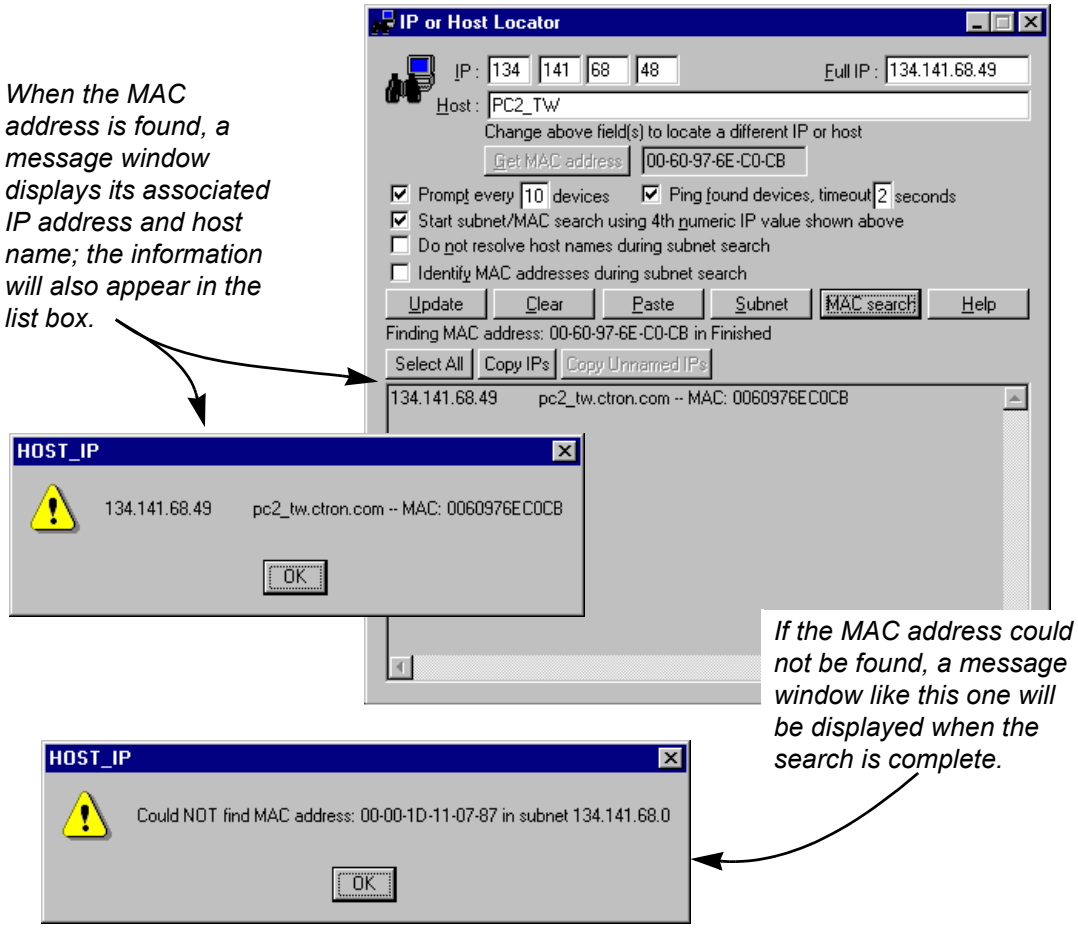
When the MAC address is found, a message window displays its associated IP address and host name; the information will also appear in the list box.
IP or Host Locator
IP : 134 141 68 48
Full IP : 134.141.68.49
Host : PC2_TW
Change above field(s) to locate a different IP or host
Get MAC address
00-60-97-6E-C0-CB
Prompt every 10 devices
Ping found devices, timeout 2 seconds
Start subnet/MAC search using 4th numeric IP value shown above
Do not resolve host names during subnet search
Identify MAC addresses during subnet search
Update
Clear
Paste
Subnet
MAC search
Help
Finding MAC address: 00-60-97-6E-C0-CB in Finished
Select All
Copy IPs
Copy Unnamed IPs
134.141.68.49 pc2_tw.ctron.com -- MAC: 0060976EC0CB
HOST_IP
134.141.68.49 pc2_tw.ctron.com -- MAC: 0060976EC0CB
OK
If the MAC address could not be found, a message window like this one will be displayed when the search is complete.
HOST_IP
Could NOT find MAC address: 00-00-1D-11-07-87 in subnet 134.141.68.0
OK

Figure 3-19. Find MAC Address Search

Using the Subnets Window

With the Subnets window, you can define and view subnet entries which can be used in Discover scripts. The Subnets window is accessed by selecting **View—>Subnets** from the NetSight Element Manager primary window menu bar. Any subnet entries that appear in the Subnets window will also appear in the Subnets group box on the **Device** page of the Discover Properties window, and vice versa (see **Defining and Selecting Subnets**, page 3-7).

The Subnets window will list subnet entries for any subnet on which an IP or Subnet Discovery has been previously performed. For example, if you have previously performed an IP Discovery using the IP Address range of 132.177.118.0 —> 132.177.118.100, the Subnets window will contain an entry for the 132.177.118.0 subnet (assuming that this is a Class B network, with a ***subnet*** mask of 255.255.255.0).

If you have manually added a device to a view or a map, a subnet entry will be listed for that device's subnet as well. For example, if you manually add a device with an IP address of 132.177.118.24 to any of your views or maps, the Subnets window will contain an entry for the 132.177.118.0 subnet (assuming that this is a Class B network, with a ***subnet*** mask of 255.255.255.0).

This subnet information is used to determine the map hierarchy created via the automatic subnet mapping feature; for more information on this feature, see ***Generating Submaps Automatically*** *in Chapter 6,* ***Creating Network Maps***

When defining subnet entries using the Subnets window, you use the same Subnet Properties window that is used in the Discover application when defining subnets for use in a Discover script.

To launch the Subnets window from the NetSight Element Manager primary window:

1. Select **View—>Subnets** from the primary window menu bar. The Subnets window, Figure 3-20, opens.

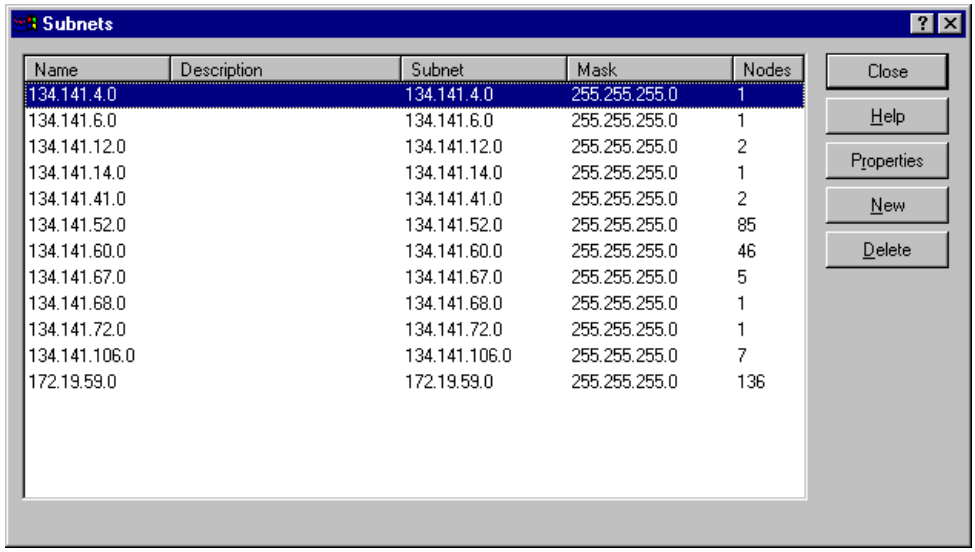

Figure 3-20. The Subnets Window

The Subnets window displays entries for all manually-defined (via this window or the Discover application) subnets; if you have manually added a node to the central node database, a subnet entry will be listed for that node's subnet as well.

Each subnet entry will be listed with its name, description, the network and subnet portions of its IP Address, its subnet mask, and count of the nodes in the central node database that reside on that subnet. A subnet's name, description, and address are all defined using the Subnet Properties window. Automatically-created subnet entries are named with their IP address by default; no default description is provided.

TIP

*Note that the Nodes value displayed here does not have the same meaning as the Nodes value displayed in the Subnet Properties window (Figure 3-6, page 3-8); that value indicates the maximum number of nodes that **could be** included in a subnet defined by the given subnet mask, and does not indicate how many nodes on a specific subnet have been added to the central node database.*

Creating and Modifying a Subnet Entry

To create a new subnet entry or modify an existing one:

1. **To define a new subnet**, click **New**; or click the **right** mouse button anywhere in the Subnets window, and select **New** from the resulting menu.

   To modify an existing subnet, highlight the subnet entry that you wish to change, and click **Properties**; or click the **right** mouse button on the subnet entry that you wish to modify, and select **Properties** from the resulting menu.

   The Subnet Properties window (Figure 3-7, page 3-10) opens.

2. Follow the steps outlined in **Defining and Selecting Subnets**, beginning on page 3-7. The procedure for defining and modifying subnet entries is the same from this window or when configuring a discover script.

3. Click **OK** to save your subnet discovery parameters and exit the Subnet Properties window. The name, description, address, and subnet mask of the defined subnet opens in the Subnets window.

Deleting a Subnet Entry

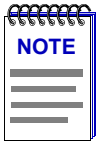

*Deleting a subnet entry **will not** delete that subnet's nodes from NetSight Element Manager's central node database.*

1. In the Subnets window, highlight the name of the subnet(s) to be deleted, and click **Delete**.

   or

   In the Subnets window, click the **right** mouse button on the subnet to be deleted, and select **Delete** from the resulting menu.

The selected subnet entry will be deleted from the Subnets window, as well as from the Subnets group box on the **Device** page of the Discover Properties window (see **Defining and Selecting Subnets**, page 3-7, for more information).

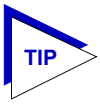

*If you delete a subnet entry from either the Subnets window or from the Subnets group box on the Device page of the Discover Properties window, that subnet will not be included in any automatic subnet mapping that takes place. See **Generating Submaps Automatically** in Chapter 6, **Creating Network Maps**, for more information on subnet mapping.*

Chapter 4

List Views

Using list views to view node information; using the List View workspace; creating and editing list views; changing the appearance of your list views; excluding nodes; adding network elements to the node database; editing node properties

---

NetSight Element Manager stores information about all discovered or manually-inserted devices in one place: its central node database. This information can be viewed in a wide variety of customizable ways via List Views, Tree Views, and Maps.

About List Views

One of the ways NetSight Element Manager allows you to see, organize, and work with the nodes in your database is through List Views. List Views, as their name implies, display the contents of your central node database in a collection of lists. These lists can be configured in a variety of ways: you can choose the information about each node that you wish to display, the order in which that information will be displayed, and the font used to display it; you can also choose which nodes you wish to display in a list by configuring a filter which selects or omits nodes based on the characteristics you choose. Specific nodes can also be excluded from individual List Views, while remaining present in the central node database. List Views provide enormous flexibility in the *display* of your central node database without altering its contents.

NetSight Element Manager provides you with several default list views, including a read-only view called the All Nodes List View, which always displays the entire contents of your central node database. Each List View can be accessed and (where appropriate) edited via the List View Workspace; the workspace also allows you to create new list views, and organize your views via folders.

The List View Workspace

Using the List View workspace you can access and edit existing views, create new ones, and organize your views by arranging them in a hierarchy of folders.

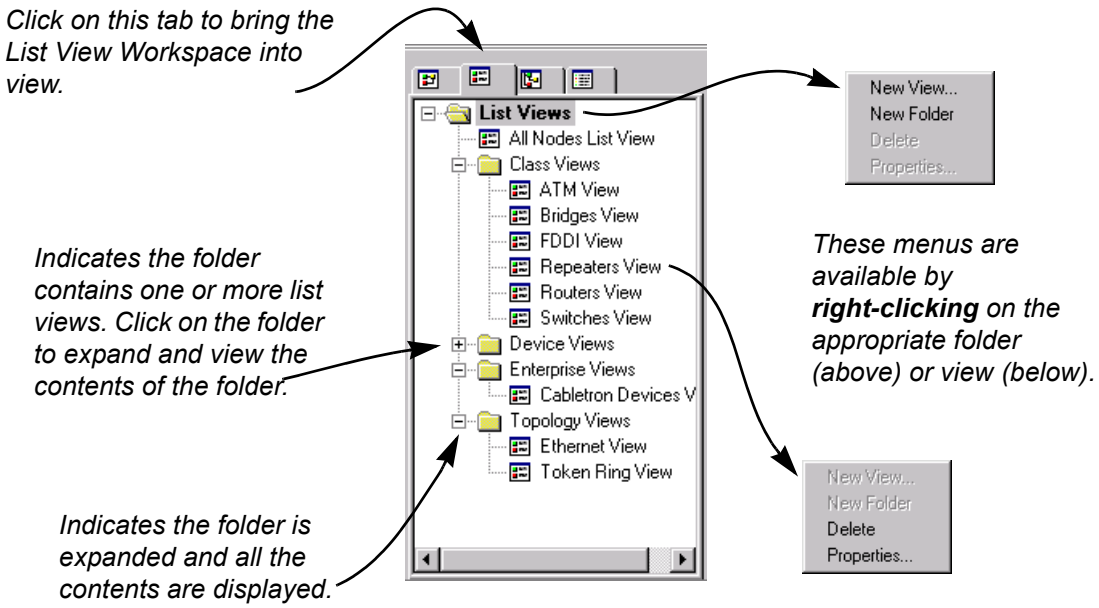

Figure 4-1. Getting Around the List View Workspace

The default views provided with NetSight Element Manager are arranged in a series of folders which group them by type. All of these default views — except the read-only All Nodes List View — can be edited or deleted; you can also add additional views to each of the default folders, or empty the folders and remove them entirely.

Adding a Folder to the Workspace

As a means of organizing a collection of List Views, you can add folders to the List View page of the workspace; each folder you create will be contained within the base List Views folder, and each can contain one or more list views, as well as additional folders. Note that folders are a workspace display convention only, and they can only be created from within the workspace itself.

To create a folder:

1. Display the List View page of the workspace by clicking on the tab in the Workspace View.
2. Click the **right** mouse button on the folder to which you want to add a new folder; that folder will be highlighted (as illustrated in Figure 4-2), and a menu opens.

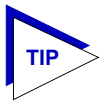

To add a folder at the top level of the List View hierarchy, click to select the main List Views folder.

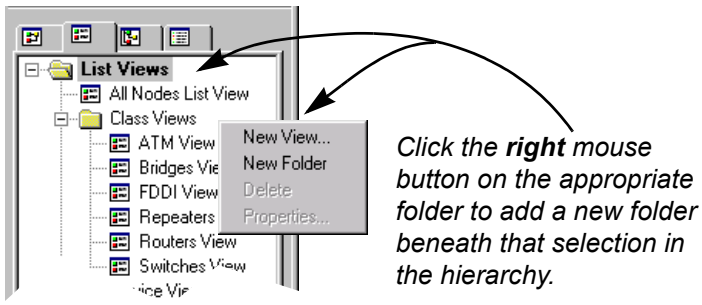

Figure 4-2. Adding a New Folder to the List View Workspace

3. In the resulting menu, click to select **New Folder**. A new folder will automatically be added in the branch you have selected.

Renaming a Folder

By default, all new folders are named with the designation **New Folder X**, where X represents a sequential index; you can rename your folders, as follows:

1. Click once to highlight the folder whose name you wish to change, then after a brief pause click again. The folder name will be enclosed in an editable text box (as illustrated in Figure 4-3), and the existing name will be highlighted.

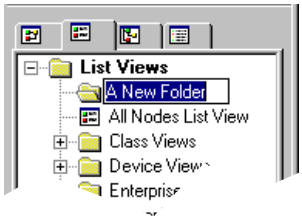

Figure 4-3. Changing a Folder Name

2. Click again to place the cursor anywhere in the name, and edit as desired; or, simply begin typing to replace the highlighted text entirely.
3. Press **Enter** to set your change.

The All Nodes List View

The All Nodes List View provides the most comprehensive look at the contents of your central node database. This special default view will always appear at the root of your List View hierarchy (as displayed in the List View workspace), and cannot be edited or deleted; by design, it displays information about every node in the central node database.

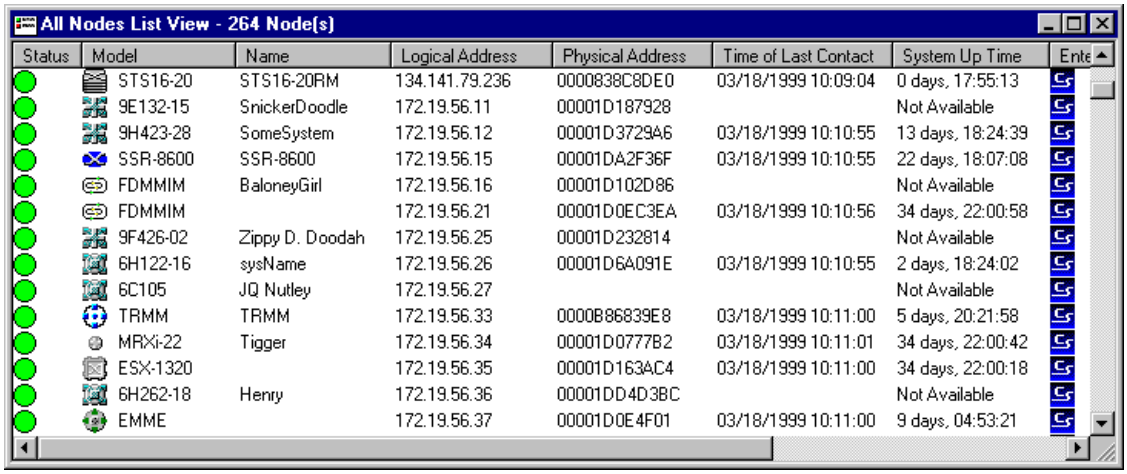

Figure 4-4. The All Nodes List View

The All Nodes List View display includes 11 of the 13 possible descriptive columns available, including Status, Model, Name, Logical Address, Physical Address, Time of Last Contact, System Up Time, Enterprise, Classes, Topologies, and Status Group; Label and Event information is omitted (but can be selected for other views). As with the other List Views, the contents of the All Nodes List View can be sorted by any one or more of the available columns, and its display format can be changed; the nodes which are displayed (and the information which is displayed about them), however, cannot.

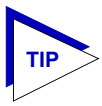

In addition to changing the arrangement of information in the All Nodes List View, the one other feature you can edit is the font used to display list view information. For more information about changing fonts on this and other list views, see ***Selecting a Display Font****, page 4-14.*

Other Default List Views

In addition to the All Nodes List View, a variety of other default views — organized in folders by type — is also provided; these views can be used as is, edited to better suit your needs, or deleted entirely. Each of these views includes a pre-set filter which excludes all nodes not of the type specified by the view's name.

Default views provided are:

Class Views

The default class views filter devices based on their functional class, displaying only those devices which provide the noted functionality: ATM (the ATM View), bridging (the Bridges View), FDDI (the FDDI View), repeating (the Repeaters View), routing (the Routers View), or switching (the Switches View). Devices which belong to more than one of these functional classes will be displayed in all appropriate views.

Device Views
A single default Device View — the All Devices View — displays all IP devices.

Enterprise Views
A single default Enterprise View allows you to display only your Cabletron and Enterasys devices.

Topology Views
Two default topology views filter out all devices that do not provide Ethernet (Ethernet View) or Token Ring (Token Ring View) topology support.

List View Display Options

In addition to selecting both the nodes and the information about those nodes that will be displayed, you can also select the arrangement of this information in a selected list view window. These display format options apply to all views, including the All Nodes List View.

Using the List View Toolbar

The appearance of the list views you create can be changed according to which information needs to have priority in the view, and how you want it displayed. The List View toolbar — and the associated options available via the primary window **View** menu — will enable you to change the size and amount of information displayed in your list view by choosing one of four display options:

Figure 4-5. The List View Toolbar

Large Icon Format	(**View—>Large Icon**) The nodes are represented by large model icons with labels below them, labeled and sorted by assigned Name (*sysName* by default). Icons are arranged in rows and sorted left to right and top to bottom.
Small Icon Format	(**View—>Small Icon**) The nodes are represented by small model icons with their labels to the right, labeled and sorted by assigned Name (*sysName* by default). Icons are arranged in rows and sorted left to right and top to bottom.

Small Icon Detail Format	(**View—>Small Icon Detail**) Each node appears as a line in a multicolumn format. The columns consist of the information you have chosen to display from the List View Properties window (or the default columns, in the case of the All Nodes List View), in the order you have specified there. For each displayed information type which includes an icon, the small (16x16 pixel) graphic will be displayed. This is the default view format; it will be applied each time NetSight Element Manager is launched.
Large Icon Detail Format	(**View—>Large Icon Detail**) Each node appears as a line in a multicolumn format. The columns consist of the information you have chosen to display from the List View Properties window (or the default columns, in the case of the All Nodes List View), in the order you have specified there. For each displayed information type which includes an icon, the large (32x32 pixel) graphic will be displayed.

Any change made to the display format is applied to the currently selected view for the duration of the current NetSight Element Manager session or until another View selection is made; once NetSight Element Manager is closed and re-opened, however, all views will revert back to the default Small Icon Detail format.

The List View toolbar also contains the Find icon:

Find	This feature allows you to search for a specific node in a List View by any of its associated information; see **Searching a List View**, page 4-7, for more information on this feature.

Sorting Nodes

By default, List Views displayed in the Small Icon Detail and Large Icon Detail formats described above are sorted in Logical (IP) Address order. However, these views can also be re-sorted by the information in one or more columns. Text fields are sorted alphabetically, numeric fields are sorted numerically, and mixed fields are sorted alpha-numerically; Enterprise, Classes, and Topology values are sorted numerically by assigned identifiers rather than alphabetically by group name.

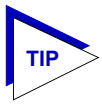

*If **Status** is selected as a sort column, the sort will update automatically as necessary.*

Sorting is accomplished by clicking on the column heading button for the information you wish to sort by. Clicking on the same column heading more than once alternately sorts the selected information in ascending and descending order; clicking on more than one column heading creates a multi-level sort. The current sort order is indicated by a number

and a greater-than or less-than sign displayed in the selected column settings: a greater-than sign (>) indicates that the information in the associated column is sorted in **ascending** order; a less-than sign (<) indicates that the information is sorted in **descending** order. You can choose as many levels as you want for your sort, and you can mix ascending and descending sorts in different columns. To return the view to the default sort order, select the **Reset Sort Order** option available on the right-mouse menu.

Figure 4-6, below, illustrates the default (Logical Address) sort order, and the same view with a multi-level sort applied. Note that the multi-level sort mixes ascending and descending orders in the different columns.

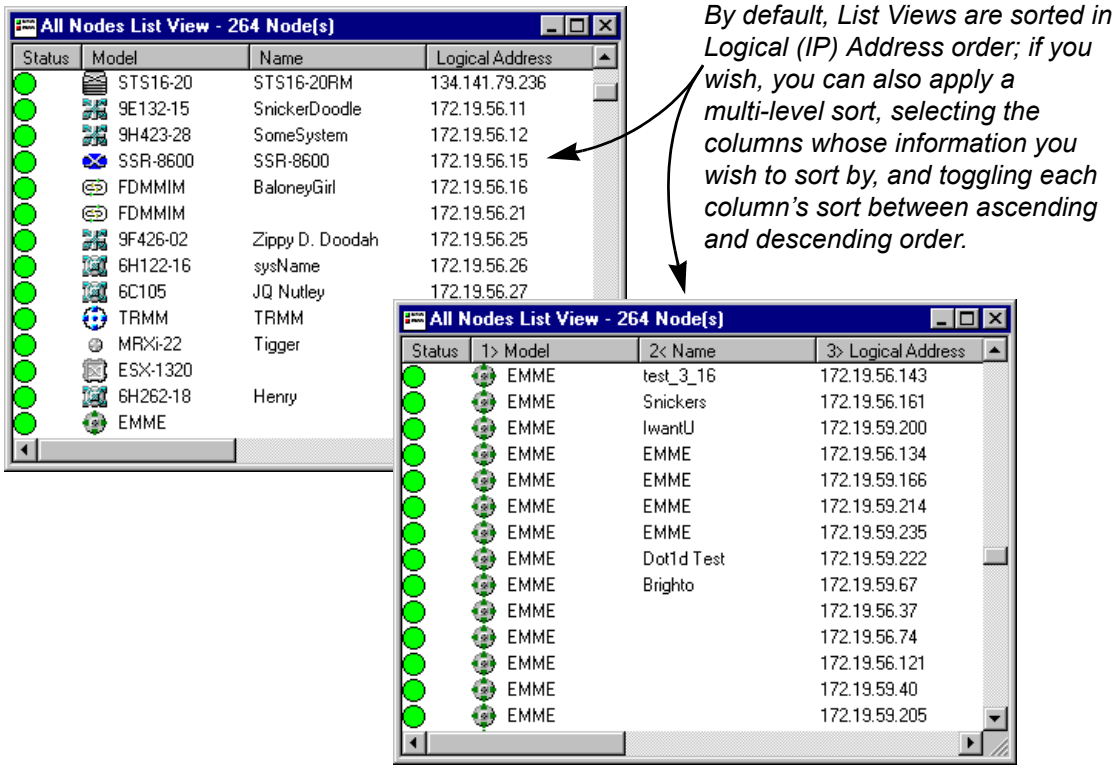

Figure 4-6. Sorting Nodes

Searching a List View

Using the Find feature, you can search a List View for any value displayed in its columns. The Find window will search the contents of the active List View for a text string (with options for searching a partial string and for case sensitivity), or for a specific value displayed in a selected column type. You can also search upwards or downwards from your current position, and choose to continue the search when the top or bottom of the displayed view is reached.

To access the Find window:

1. Open or activate the List View that you wish to search.
2. Select **Edit—>Find** from the primary window menu bar.

   or

   Click on the List View toolbar.

The List View Find window, Figure 4-7, opens.

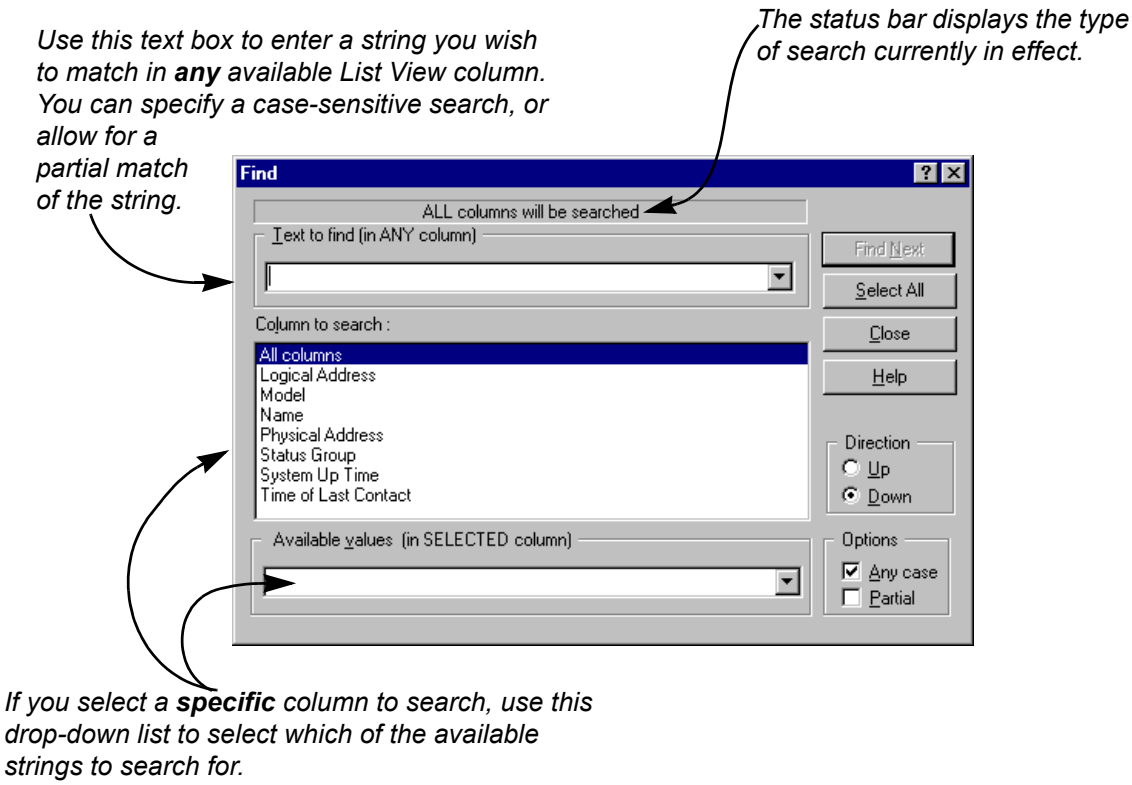
Use this text box to enter a string you wish to match in **any** available List View column. You can specify a case-sensitive search, or allow for a partial match of the string.
The status bar displays the type of search currently in effect.
Find
ALL columns will be searched
Text to find (in ANY column)
Find Next
Select All
Close
Help
Column to search :
All columns
Logical Address
Model
Name
Physical Address
Status Group
System Up Time
Time of Last Contact
Direction
Up
Down
Available values (in SELECTED column)
Options
Any case
Partial
If you select a **specific** column to search, use this drop-down list to select which of the available strings to search for.

Figure 4-7. List View Find Window

You can perform two different kind of searches with this window: you can enter any text string you wish to match fully or partially in any column (using the top text box); or you can search for any one of the available values in a selected column (using the lower text box).

To search for a value in ANY column:

1. In the **Column to search** list box, make sure the **All columns** option is selected.

2. In the **Text to find (in ANY column)** text box at the top of the window, enter the string you wish to match. Note that this field will not be available unless the **All columns** option is selected in the **Column to search** list box. For your convenience, the ten most recent search strings will be listed in a drop-down menu.
3. To make your search case-sensitive, click to deselect the **Any case** option in the lower right corner of the window. To match any case, leave the option selected.
4. If you want your search string to be matched as a whole word, click to deselect the **Partial** option. If you want to allow a partial match, leave the option selected.
5. Use the **Direction** options to specify whether to search **Up** or **Down** from the entry currently selected in the active List View. Whichever direction you choose, you will be given the opportunity to start over when the top or bottom of the list is reached.
6. Click on **Find Next** to begin the search. Each match will be highlighted as it is found; to find additional matches, click on **Find Next** again.

   or

   Click on **Select All** to immediately highlight all matches.

If the selected value is not found, a message window will notify you of the failure.

To search for a value in a SELECTED column:

1. In the **Column to search** list box, click to select the column whose values you wish to search. You can restrict your search to any one of the following columns, as long as it is displayed in the active List View: **Logical Address**, **Model**, **Name**, **Physical Address**, **Status Group**, **System UpTime**, and **Time of Last Contact**.
2. Click on the arrow to the right of the **Available values (in SELECTED column)** field to select the value you wish to search for. All values which appear in the selected column will be available for your search; these are the *only* values that can be used for this kind of search. Note that this field will *not* be available if the **All columns** option is selected in the **Column to search** list box.
3. Use the **Direction** options to specify whether to search **Up** or **Down** from the entry currently selected in the active List View. Whichever direction you choose, you will be given the opportunity to start over when the top or bottom of the list is reached.
4. Click on **Find Next** to begin the search. Each match will be highlighted as it is found; to find additional matches, click on **Find Next** again.

   or

   Click on **Select All** to immediately highlight all matches.

Creating and Editing List Views

You can create a variety of list views customized to meet your network management needs, either by editing any of the default views provided (except the All Nodes List View) or by creating entirely new views. User-defined filter criteria allow you to select which *nodes* will be displayed; List View display properties allow you to select which *information* about those nodes will be displayed.

To create a new list view:

1. Select **Window —> New List View** from the primary window menu bar.

   or

   Click on the toolbar.

   or

   In the List View Workspace, click the *right* mouse button on any folder or in any blank area, and select **New View** from the menu that appears.

   The New List View window, Figure 4-8, opens.

*The **Columns** check list allows you to choose which attributes are displayed in the list view; you can also change the order in which this information will be displayed. For more information, see* ***Selecting Node Information to Display**, page 4-11*.

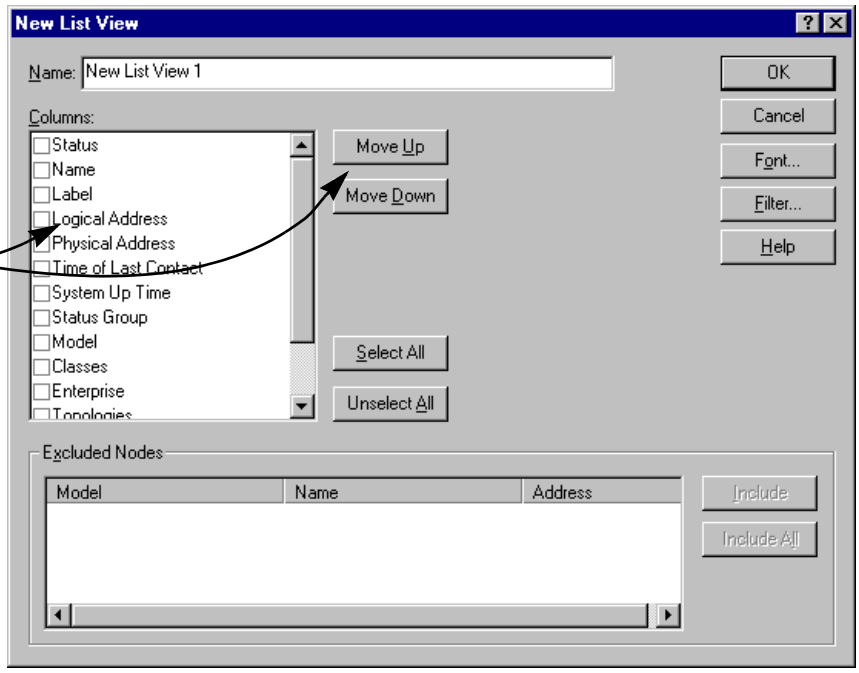

Figure 4-8. The New List View Window

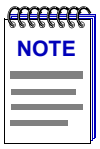

When you add a new list view by any of the means described above, it will be added to the folder which is currently selected in the List View workspace; if you add a list view by right-clicking in a blank area of the workspace, the new view will be added to the root folder. Once a List View has been added, it cannot be moved from folder to folder. All folders and views are listed in the workspace in alphabetical order.

To edit an existing list view:

1. With the appropriate list view open and selected in the NetSight Element Manager primary window, select **View —> Properties** from the primary window menu bar.

   or

   In the List View Workspace area, click the *right* mouse button on the name or icon associated with the view you wish to edit, and select **Properties** from the resulting menu.

   or

   Click with the *right* mouse button anywhere in the view you wish to edit, and select **List View Properties** from the resulting menu.

   The List View Properties window will be displayed.

The procedures for creating a new List View and editing an existing one are virtually identical; they are described in the following sections.

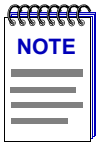

If you open the List View Properties window for the All Nodes List View, note that all fields are grayed out; remember, the only thing you can change about this view is the font in which its information is displayed. See ***Selecting a Display Font****, page 4-14, for more information.*

Selecting Node Information to Display

The main List View Properties window allows you to assign or change the name of your list view, select the information that will be displayed for each node, and select the order in which it will be displayed.

To set up or modify a list view:

1. By default, a new list view will be named **New List View X**, where X will increment sequentially as necessary to provide a unique name. To change this default (or the original name assigned to any view), highlight the text in the **Name** field and delete, then enter a new name.

2. In the **Columns** list box, click to select the information you'd like displayed about each node in the view; click again to deselect any option you do not wish to view. Each selection in this list will add a column of information to your list view. Each of the available options is described below; you must select at least one column to create or save your view.

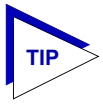

*You can also use the **Select All** and **Unselect All** buttons to select or deselect all available column options.*

3. The column options you have selected display in the list view in the same order they are listed in the **Columns** field; to change this order, click to select the option you want to move, and use the **Move Up** and **Move Down** buttons to change its position as necessary. Repeat to order all selected options as desired.

A list view configured in this way will display the information you have selected for *all* nodes in the central node database, using the *default* display font. To change the node types that will be included in your view, use the **Filters** option (described in **Selecting the Nodes to Display: Using Filters**, page 4-15); to change the display font, use the **Font** option (described in **Selecting a Display Font**, page 4-14).

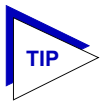

*For more information about the view-specific Excluded Nodes list displayed at the bottom of the List View Properties window, see **Excluding Nodes from an Individual View**, page 4-18.*

Node Information Options

List views can include up to 13 columns of information about the nodes they display. Information available for display includes:

Status

Indicates the status of NetSight Element Manager's contact with the node. The possible states are:

- **Critical** (Red) - The node is in critical condition; it is no longer in contact with NetSight Element Manager.
- **Warning** (Yellow) - The status icon turns yellow to indicate that NetSight Element Manager has received a trap from the represented node; the icon will remain yellow until the trap is acknowledged.

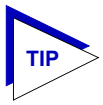

If any node in your central node database is in a Critical or Warning state, this status condition will be propagated right up to your desktop via the minimized NetSight Element Manager icon — the Task Bar icon representing NetSight Element Manager will change from the inactive window title bar color to the active window title bar color.

- **Unknown** (Blue) - A blue status icon indicates that NetSight Element Manager has not yet been able to contact the represented node, and has not yet determined whether or not that contact will be successful. A node icon is blue only temporarily, while NetSight Element Manager determines whether or not contact can be made: if contact is successfully made, the icon will turn green; if not, it will turn red. Performing an individual node status query (using the **Manage—>Query Status** option from the primary window menu bar, or the **Query Status** option available from the right-mouse menu) will temporarily turn a node status icon blue, while the status query is being made.
- **Normal** (Green) - NetSight Element Manager is successfully communicating with the represented node, and the node is operating properly on the network. Green is also the default color for a map icon which is not currently associated with an IP address.
- **Undefined** (Gray) - A gray status icon for a node indicates that polling has been disabled, and the node's current status is not known.

Name

This column displays the user-defined name of the node. You can assign a name for the node in the **General** properties page in the Device Properties window. By default, devices will display any configured MIB II *sysName* value.

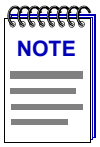

*Note that changing a node's name via its Properties window changes **only** the information in NetSight Element Manager's central node database; no name changes are set at the device. If you wish to leave the MIB II* sysName *displayed, **and** assign an additional label, use the Label field described below. For more information on configuring node properties — including setting both Name and Label values — see* ***Editing Device Properties**, page 4-29*.

Label

Displays a user-set alias for a node; no default label is provided.

Logical Address

Displays the IP (Internet Protocol) address for devices.

Physical Address

Displays the MAC address for devices.

Time of Last Contact

Lists the date and time of NetSight Element Manager's last successful contact with the node.

System Up Time

The length of time the device has been operating since its last boot up, displayed in a days, HH:MM:SS format.

Status Group
If the device is part of a Distributed LAN Management status group, it will be indicated here. Both the name assigned to the DLM group and the node's position in that group ("DLM Agent" or "DLM Client of (IP Address)") will be displayed. For more information on configuring and using Distributed LAN Management functions, see the **DLM Status Groups** chapter in the ***NetSight Element Manager Tools Guide***, included with your document set.

Model
Indicates the model name of the device, along with the icon which represents that model type.

Classes
Displays the functional classes supported by the selected node (Router, Bridge, Switch, ATM, etc.), represented by the appropriate class icon(s).

Enterprise
Displays the enterprise, or vendor — Cabletron, Enterasys, etc. — which manufactured the selected node.

Topologies
Displays the topology (or topologies) supported by the node.

Events
Indicates whether or not the Alarm and Event log has received traps and/or events related to the node.

Selecting a Display Font

The Font Window allows you to select the font style and size for any displayed text. You can change the font setting for *any* list view, including the otherwise-uneditable All Nodes List View.

1. Click on **Font** from the appropriate New List View or List View Properties window. The Font window, Figure 4-9, opens.

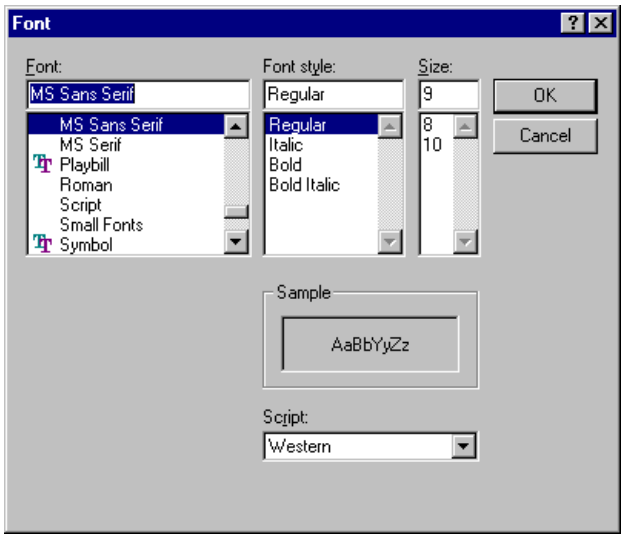

Figure 4-9. Font Window

Change your font settings as desired using the **Fonts**, **Font Style**, and **Size** fields, then click on **OK** to set your changes.

Selecting the Nodes to Display: Using Filters

Using the View Filter window, you can set up filters that determine which of the nodes stored in the central node database will be displayed in your list view. You can filter nodes based on specific **Subnet**, **Status Group**, **Enterprise**, **Class**, and/or **Topology** attributes.

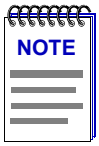

With the exception of the All Nodes List View, the default views provided by NetSight Element Manager were created with the use of filters. For example, the Cabletron Devices View filters out all devices except Cabletron and Enterasys devices.

Remember, you can edit any of the default views except the All Nodes List view; the default views can also serve as models to help you create your own custom views.

To open the View Filter window:

1. Click on **Filter** from the New List View or List View Properties window. The View Filter window, Figure 4-10, opens.

The check boxes above the filter categories allow you to either select or de-select all of the specific attributes listed. You must leave at least one attribute selected in each category.

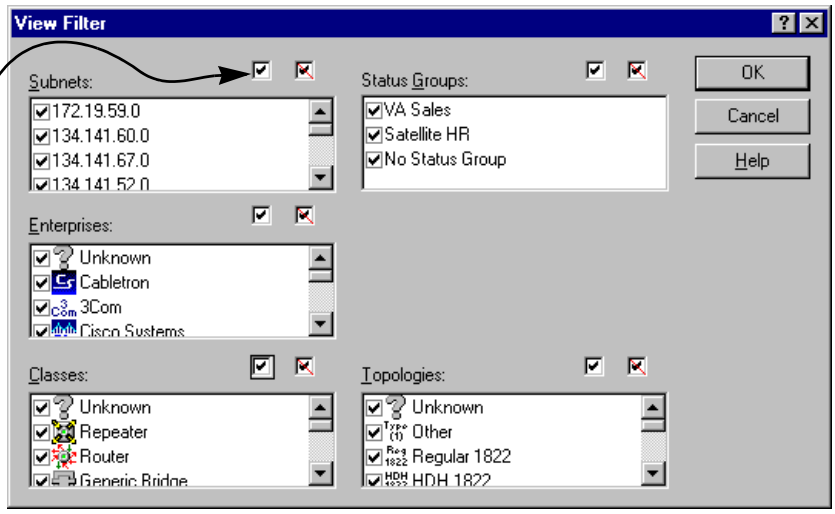

Figure 4-10. View Filter Window

Filter Categories

Available filter categories are:

Subnets

NetSight Element Manager automatically creates a subnet entry each time a node is added to the database; subnet entries can also be defined via the Discover application and/or via the **View—>Subnets** option in the primary window menu bar. All subnets which have been created will be displayed in this list box; use this list to select the subnets whose nodes you wish to display in your list view.

TIP

*See Chapter 3, **Discovering Nodes**, for more information on creating and using subnet entries.*

Status Groups

Any DLM Status Groups you have created will be displayed in this list box. You can choose to display only those nodes which are part of one or more selected status groups, or those which are part of *any* defined status group.

Enterprises

The Enterprise category allows you to filter nodes by hardware or software manufacturer. The vendor codes available for filtering are Cabletron, 3Com, Cisco Systems, Novell, Silicon Graphics, IBM, Sun Microsystems, Hewlett-Packard, Microsoft, Xyplex, Cayman, Digital, Compaq, and Enterasys. You can also choose to filter on nodes whose vendor code is unknown.

Classes

Lists the functional classes that NetSight Element Manager can detect. Classes available for filtering include:

- Repeater
- Router
- Generic Bridge
- Router/Gateway
- Ethernet/Token Ring Bridge
- Switch
- Ethernet/FDDI Bridge
- Ethernet/WAN Bridge
- PBX
- Multi-Topology Bridge
- FDDI Concentrator
- Terminal Server
- Server
- ATM

Nodes which support more than one functional class will be displayed in any view for which one of its classes is selected. You can also choose to filter on nodes whose functional class is unknown.

Topologies

Lists the topologies that NetSight Element Manager will recognize. Topologies listed include:

- Regular 1822
- HDH 1822
- Local Talk
- RFC 877 X.25
- Ethernet CSMACD
- ISO 8802.3
- ISO 8802.6 MAN
- ISO 88022 LLC
- 3 Mbit Ethernet
- StarLAN
- Token Bus
- Token Ring
- HyperChannel
- FDDI
- Lap B
- T1
- T-3
- E1
- Basic ISDN
- Primary ISDN
- PPP
- Serial Line IP
- CLNP/IP
- XNS/IP
- Software Loopback
- Ultra Tech
- SMDS
- Proteon 80Mbit
- Proteon 10Mbit
- ATM
- AAL5 over ATM
- RS-232
- Parallel Port
- Arcnet
- Arcnet Plus
- SMDS DXI
- SONET
- SONET Path
- SONET VT
- SDLC
- V.35
- Frame Relay
- Frame Relay Service
- HIPPI
- HSSI
- X.25
- Proprietary Multiplexing
- X25PLE
- M10X25
- Proprietary Virtual/Internal
- Generic Modem
- Proprietary PPP Serial
- SMDS Intercarrier Interface
- Other

As with classes, nodes which support more than one topology will be displayed in any view for which one of its topologies is selected. You can also choose to filter on nodes whose topology is unknown.

Creating the Filter

When you create a filter, you select the specific attributes in each category that define the nodes you *want* to display, and de-select the attributes that define nodes you do *not* want to display. Filter categories are linked by logical ANDs, and attributes within categories are linked by logical ORs; this means that, in order to be displayed in a filtered view, a node must match *at least one* selected attribute in *each* category. You must leave at least one attribute selected in each category.

To specify a view filter:

1. In each filter category list box, click on the check box next to the specific attribute to include or exclude associated nodes from the list view.

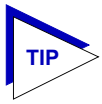

You can also use the select and deselect check boxes located above each filter category list box to either select or deselect all of its specific attributes; remember, however, that you must leave at least one attribute selected in each category in order to create your filter.

2. Click on **OK** to set the filter parameters.

Once you have finished configuring all of the parameters for your list view, click **OK** in the New List View window or List View Properties window to create the new list view or modify an existing view.

Excluded Nodes

NetSight Element Manager's architecture — which stores all node information in a central database, and lets you choose how to view it — provides you with great flexibility in the management and storage of node information. One of the ways you can manage your collected node information is through the use of Excluded Nodes. There are two levels of node exclusion: exclusion from an individual view, and exclusion from all views.

Excluding Nodes from an Individual View

Excluding a node from an individual List, Tree, or Map View removes that node from the selected view only — the node's information is still stored in the central node database, and it will be displayed as appropriate in other views. For a map view, excluding a node from the current (or active) view simply removes the icon from the map; no record of that node's former place on the map is maintained. For a List View or Tree View, however, excluding a node from the active view places it on a special excluded nodes list associated with the selected view. This list is displayed in the view's Properties window, as

illustrated in Figure 4-11. Nodes excluded from an individual view will still appear as appropriate on all other views, and their placement on maps will not be affected; they can also be re-included in the view at any time.

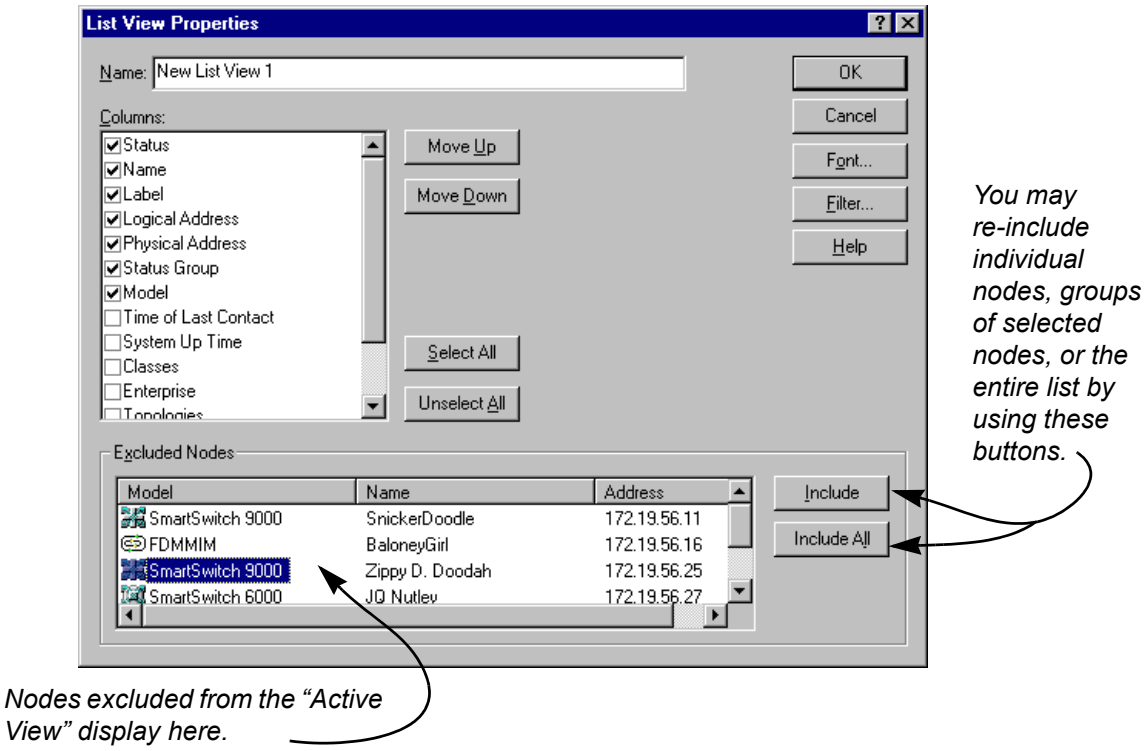

Figure 4-11. The Individual View Excluded Nodes List

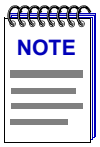

*For more information about Tree Views, see Chapter 5, **Tree Views**; for more information about maps, see Chapter 6, **Creating Network Maps**.*

To Exclude a Node from a View

To exclude a node from a specific view only:

1. Open or activate the view which contains the nodes you wish to exclude, and click (or use *shift-click* or *ctrl-click*) to select the appropriate node entries.
2. Select **Edit —>Delete** from the primary window menu bar.

   or

Select the Delete key on your keyboard.

or

Select the Delete toolbar icon.

or

Click the *right* mouse button on any of the selected entries, and select **Delete** from the resulting menu.

The Node Delete window (Figure 4-12) opens, with the node(s) you have selected for deletion highlighted in the central window.

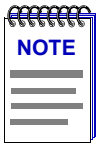

*Using **Edit —> Cut** instead of **Edit —> Delete** will remove nodes from the active view without bringing up the Node Delete window, and place them on that list view's excluded nodes list.*

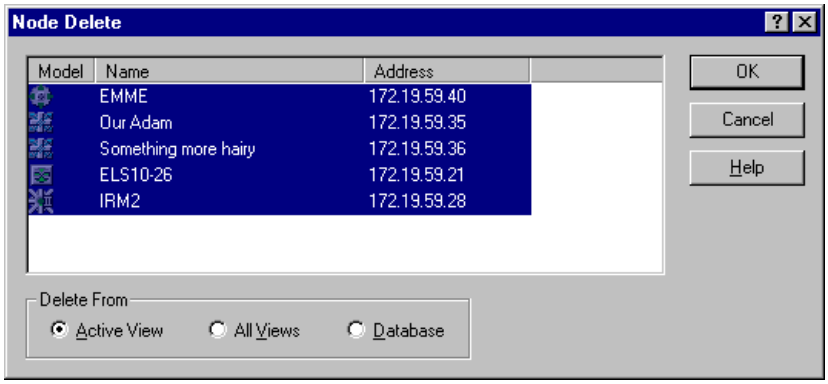

Figure 4-12. The Node Delete Window

3. Make sure the **Active View** option is selected (this option is selected by default), then click **OK** to confirm the deletion. The selected nodes will no longer be displayed in the selected view window, and will be listed in the view's properties window Excluded Nodes list.

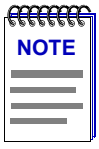

*When excluding a node from the All Nodes list and tree views, note that the "Active View" selection is not available; excluding a node from either of these views automatically excludes it from all views. See **Excluding Nodes from the Central Node Database**, page 4-22, for more information.*

To Re-include a Node in a View

To remove a node from a view's Excluded Nodes list and re-include it in the view:

1. Open or activate the view whose excluded nodes you wish to re-include.
2. Open the view's Properties window by clicking the **right** mouse anywhere in the view, and select List View Properties (or Tree View Properties) from the resulting menu.

   or

   With the appropriate list view open and active, select **View—>Properties** from the primary window menu bar.

   or

   In the List View Workspace area, click the *right* mouse button on the name or icon associated with the view you wish to edit, and select **Properties** from the resulting menu.

   The appropriate Properties window, as illustrated in Figure 4-13, opens.

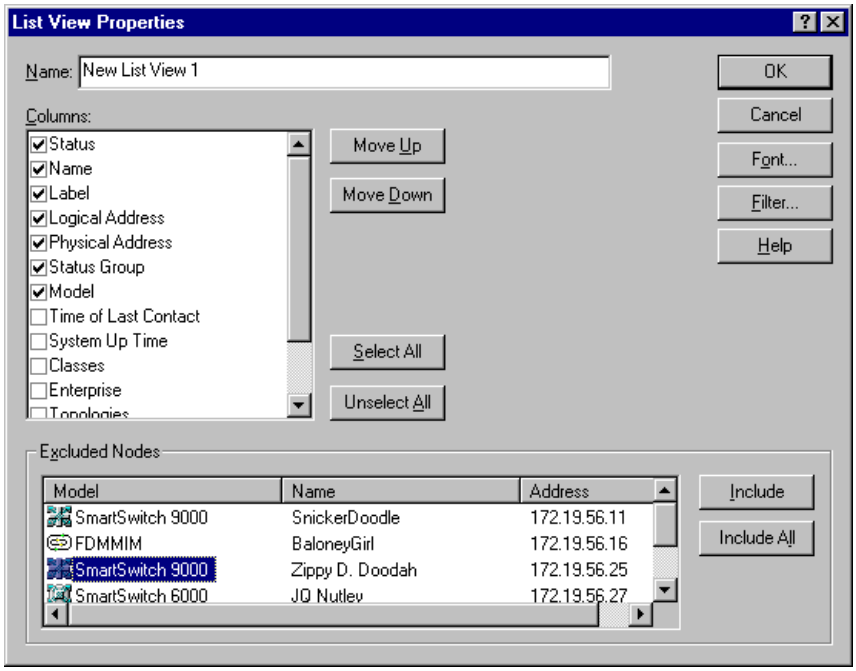

Figure 4-13. Re-including Nodes Excluded from a View

3. In the Properties window Excluded Nodes list, click to select the node or nodes you wish to re-include in the view, then click on **Include**; or, click on **Include All** to re-include all currently-excluded nodes.

4. In the resulting confirmation window, click **Yes** to confirm the inclusion, or **No** to cancel.

5. Click **OK** to close the Properties window; note that the formerly excluded nodes are again displayed in the view.

Excluding Nodes from the Central Node Database

In addition to its central node database, NetSight Element Manager maintains a second database called the Excluded Nodes database. The Excluded Nodes database serves as a kind of temporary holding area for nodes which have been excluded from all views but not yet deleted from the database. Nodes which have been placed in the Excluded Nodes database are not polled by NetSight Element Manager, and are no longer displayed in any List, Tree, or Map View; they can only be viewed via the Excluded Nodes window.

To access the Excluded Nodes window:

1. Select **View —>Excluded Nodes** from the primary window menu bar. The Excluded Nodes window, Figure 4-14, opens.

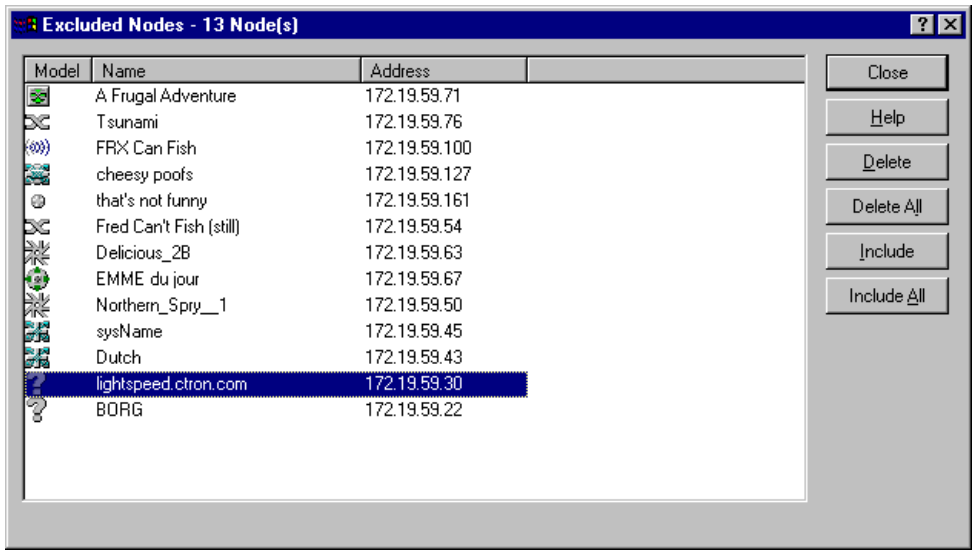

Figure 4-14. The Excluded Nodes Window

The Excluded Nodes window allows you to both view and sort the list of nodes which have been excluded from all views; from this window, excluded nodes can be returned to the central node database (and re-displayed in all appropriate views), or removed from NetSight Element Manager entirely. The excluded nodes list can also be sorted by clicking on the column heads; note that only limited information about excluded nodes is displayed.

The Excluded Nodes List and Discover

The Excluded Nodes list can be used to prevent certain nodes from being re-discovered and re-added to the central node database. Nodes which have been excluded from all views are no longer polled byNetSight Element Manager or displayed in any of its views, but they won't be discovered by the Discover application, either. If your discover scripts pick up devices you do not wish to monitor, you can simply exclude them from all views and leave them in the Excluded Nodes database indefinitely, maintaining a record of their presence on your network without the bandwidth cost of continuous polling. Nodes in the Excluded Nodes database can also be re-added to the central node database at any time, should the need arise.

For more information about the Discover application, see Chapter 3, **Discovering Nodes**.

To Exclude a Node from the Central Node Database

To exclude a node from the central node database (and add it to the Excluded Nodes database):

1. Open or activate any view (including a map view) which contains the nodes you wish to exclude, and click (or use *shift-click* or *ctrl-click*) to select the appropriate node entries.
2. Select **Edit —>Delete** from the primary window menu bar.

   or

   Select the Delete key on your keyboard.

   or

   Select the Delete toolbar icon.

   or

   Click the **right** mouse button on any of the selected entries, and select **Delete** from the resulting menu.

   The Node Delete window (Figure 4-15) opens, with the node(s) you have selected for deletion highlighted in the central window.

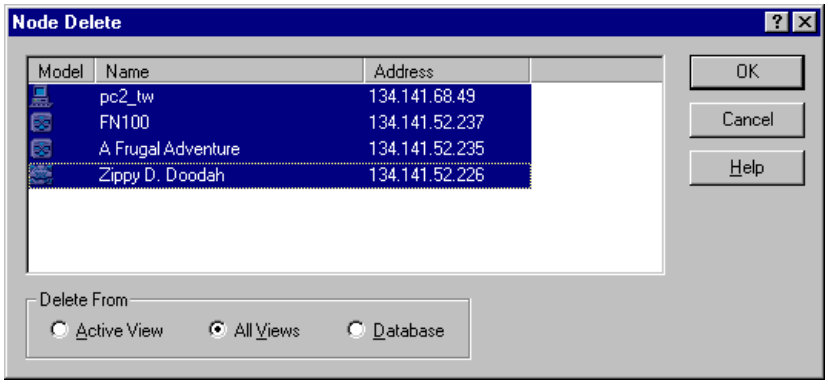

Figure 4-15. The Node Delete Window

3. Click to select the **All Views** option, then click **OK** to confirm the deletion. The selected nodes will be removed from the central node database and all views, and placed in the Excluded Nodes database.

Remember, the excluded nodes database can be viewed via the Excluded Nodes window (Figure 4-14, page 4-22), available from the **View** menu in the primary window menu bar. Excluded nodes can also be either re-included in the database, or deleted entirely.

To Re-include a Node into the Central Node Database

To move an excluded node back into the central node database:

1. Select **View —>Excluded Nodes** from the primary window menu bar to display the Excluded Nodes window.
2. In the main portion of the window, click (or use *shift-click* or *ctrl-click*) to select the nodes you'd like to re-include, then select **Include**; or select **Include All** to re-include all nodes currently in the excluded nodes list. A text box opens (Figure 4-16), asking you to confirm the request to include the selected item(s).

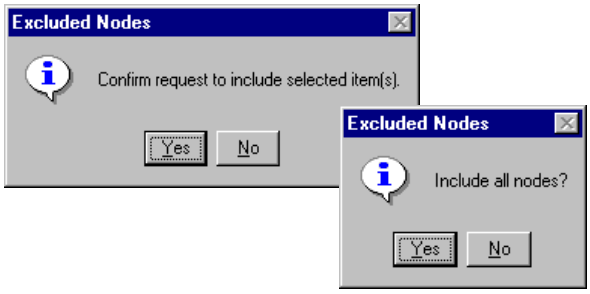

Figure 4-16. Include Confirmation Boxes

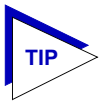

Remember, you can also click on any column header in the Excluded Nodes list to re-sort the nodes according to the information in the selected column; entries are always sorted in ascending order.

3. Click **Yes** to confirm. The node or nodes will be moved from the Excluded Nodes database to the central node database, and appear again in all appropriate list views.
4. Click **Close** to close the excluded nodes window.

Re-included nodes will be displayed as appropriate in all list and tree views, and can be re-added to map views, as desired.

To Delete a Node from the Database

You can also choose to remove nodes entirely from the NetSight Element Manager database, either directly from the Node Delete window (bypassing the excluded nodes database), or from the Excluded Nodes database itself. Nodes which have been entirely removed are not protected from re-discover, and will be automatically re-added to the database if they are discovered. Deletion from the database cannot be undone, except by re-discovering or re-inserting the nodes you want.

To delete a single node or group of nodes directly from a view:

1. Open or activate any view (including a map view) which contains the nodes you wish to delete, and click (or use *shift-click* or *ctrl-click*) to select the appropriate node entries.
2. Select **Edit —>Delete** from the primary window menu bar.

   or

   Select the Delete key on your keyboard.

   or

   Select the Delete toolbar icon.

   or

   Click the **right** mouse button on any of the selected entries, and select **Delete** from the resulting menu.

   The Node Delete window (Figure 4-17) opens, with the node(s) you have selected for deletion highlighted in the central window.

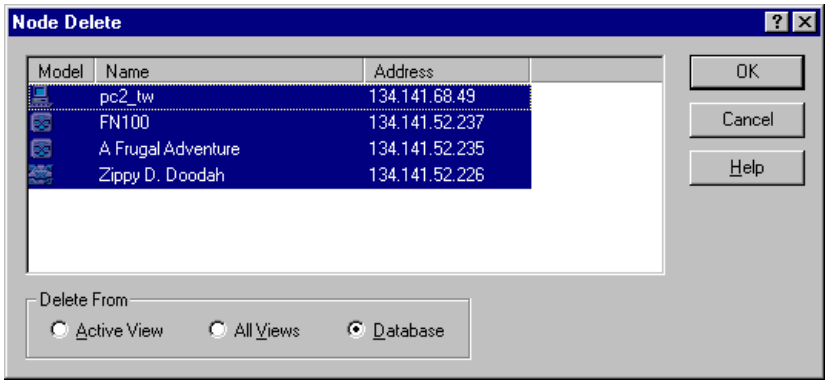

Figure 4-17. The Node Delete Window

3. Click to select the **Database** option, then click **OK** to confirm the deletion. The selected nodes will be removed from the central node database entirely.

To delete a single node or group of nodes from the Excluded Nodes window:

1. Select **View—>Excluded Nodes** from the primary window menu bar to display the Excluded Nodes window.
2. In the main portion of the window, click (or use *shift-click* or *ctrl-click*) to select the nodes you'd like to re-include, then select **Delete**; or select **Delete All** to permanently remove all nodes currently in the excluded nodes list. A text box opens, asking you to confirm the request to delete the selected item(s).

TIP

Remember, you can also click on any column header in the Excluded Nodes list to re-sort the nodes according to the information in the selected column; click again to reverse the sort order.

3. Click **Yes** to confirm. The node or nodes will be permanently deleted from the database.
4. Click **Close** to close the excluded nodes window.

Adding Network Elements to the Node Database

When you perform a network Discover operation, all discovered devices are automatically added to the central node database, and they will automatically appear in the appropriate list views. However, you can also add individual devices to the central node database via the **Insert—>Device** options available on the primary window menu bar or via the Insert Device icon on the toolbar. As with discover, nodes individually added to the central node database are automatically displayed in all appropriate List Views.

The following sections describe how to add individual devices to the central node database, and how to edit the display properties of any node.

Adding a Device

In order to add a device to the central node database, you must have a List View or Tree View open and active in the primary window. If a map view is active, an inserted device must be placed in the map before it is added to the database; if no view is open, the insert device options will not be available. For more information on maps, see Chapter 6, **Creating Network Maps**.

To add a new device to the node database:

1. In the primary window menu bar, select **Insert—>Device**.

   or

   In the toolbar, click the **Insert Device** icon.

   or

   Anywhere within an active view, **right-click** and select **New Device**, from the resulting menu.

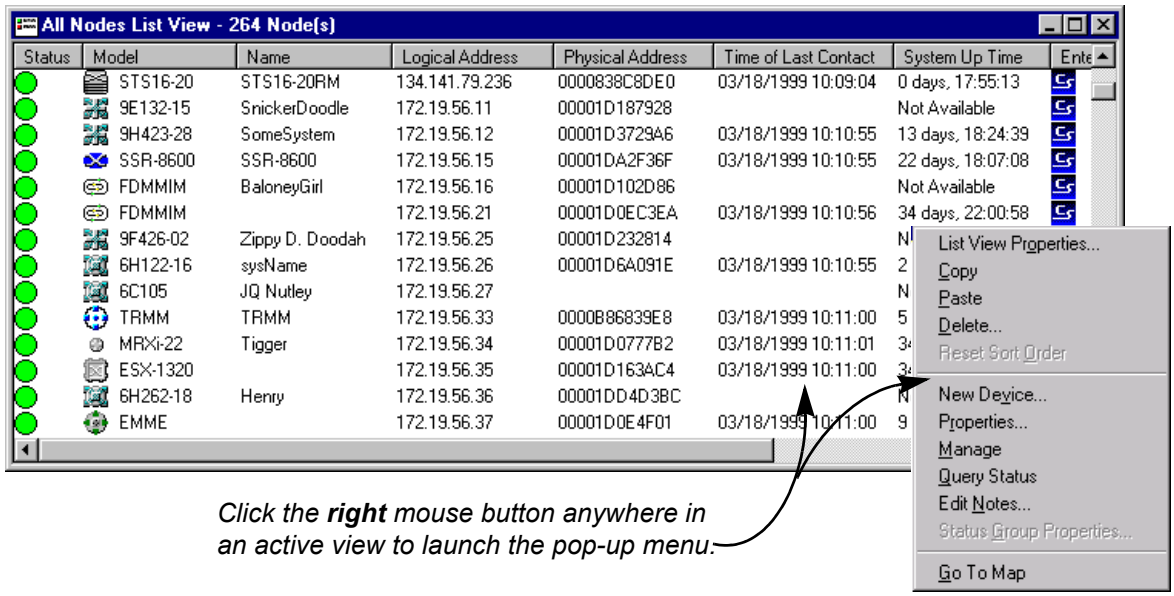
All Nodes List View - 264 Node(s)
Status
Model
Name
Logical Address
Physical Address
Time of Last Contact
System Up Time
STS16-20 STS16-20RM 134.141.79.236 0000838C8DE0 03/18/1999 10:09:04 0 days, 17:55:13
9E132-15 SnickerDoodle 172.19.56.11 00001D187928 Not Available
9H423-28 SomeSystem 172.19.56.12 00001D3729A6 03/18/1999 10:10:55 13 days, 18:24:39
SSR-8600 SSR-8600 172.19.56.15 00001DA2F36F 03/18/1999 10:10:55 22 days, 18:07:08
FDMMIM BaloneyGirl 172.19.56.16 00001D102D86 Not Available
FDMMIM 172.19.56.21 00001D0EC3EA 03/18/1999 10:10:56 34 days, 22:00:58
9F426-02 Zippy D. Doodah 172.19.56.25 00001D232814
6H122-16 sysName 172.19.56.26 00001D6A091E 03/18/1999 10:10:55
6C105 JQ Nutley 172.19.56.27
TRMM TRMM 172.19.56.33 0000B86839E8 03/18/1999 10:11:00
MRXi-22 Tigger 172.19.56.34 00001D0777B2 03/18/1999 10:11:01
ESX-1320 172.19.56.35 00001D163AC4 03/18/1999 10:11:00
6H262-18 Henry 172.19.56.36 00001DD4D3BC
EMME 172.19.56.37 00001D0E4F01
List View Properties...
Copy
Paste
Delete...
Reset Sort Order
New Device...
Properties...
Manage
Query Status
Edit Notes...
Status Group Properties...
Go To Map
Click the right mouse button anywhere in an active view to launch the pop-up menu.

Figure 4-18. Displaying the Pop-up Menu

The Insert Device window, Figure 4-19, opens.

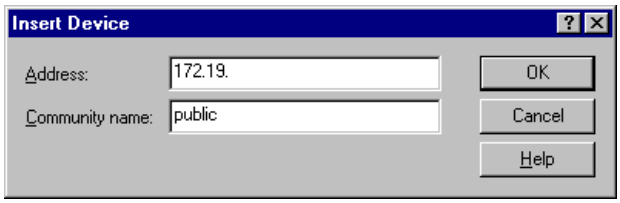

Figure 4-19. The Insert Device Window

2. In the **Address** field, enter the IP address assigned to the device you wish to add.
3. In the **Community Name** field, enter a community name that will provide at least Read access to the device you are adding. For full management access, use a community name that will provide superuser access.

TIP

*If you have configured and activated the Autofill options (either via the Welcome Wizard or via the **Tools—>Options** window), the autofill values you have set for IP address and community name display automatically in the appropriate fields. You can keep these defaults or replace them, as necessary.*

4. Click **OK** to start the insertion process. **If the device is successfully contacted**, the Device Properties window, Figure 4-21, opens. Use the pages in this window to change or configure the device's display properties (as described in the following sections), then click on **OK**.

   If the device cannot be contacted, the Insert Device window will expand to display the **Timeout** and **Retry** counters, and a message window opens (as illustrated in Figure 4-20).

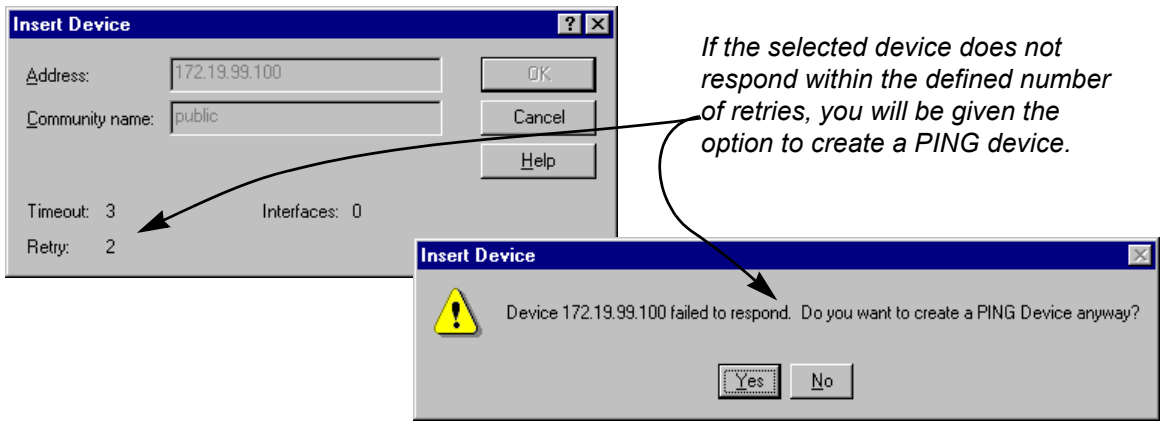

Figure 4-20. Device Timeout Windows

If you click **Yes**, the Device Properties window opens, and you can add an **Unnamed PING device** to the central node database. If you click **No**, no new device will be added.

TIP

If a map window was active when you selected the menu or toolbar option for adding a device, closing the Properties window will change the arrow pointer into a crosshair cursor. You can click this cursor in any open map window to add the appropriate device icon to the map; adding a device to a map automatically adds it to the central node database, and it displays on all appropriate list views.

Editing Device Properties

The Device Properties window, Figure 4-21, provides four tabbed pages which allow you to change and/or configure certain display and polling properties associated with the device.

Click on the tabs to display the four different pages available in the Device Properties window

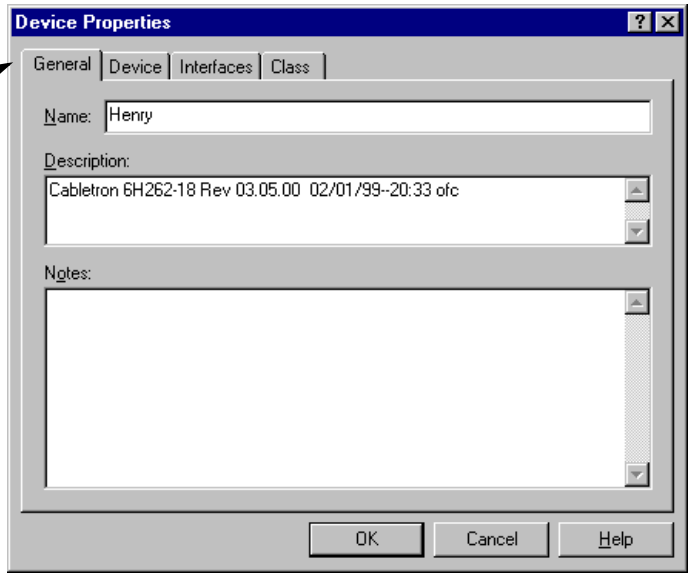

Figure 4-21. The Device Properties Window

The Device Properties window appears automatically each time a device is manually added to the node database (or specifically to a map window). To access the properties window for a device which as already been added to the database (or to a map window):

1. Open and/or activate the list, tree, or map view which contains the device whose display properties you wish to edit, and click on the device icon to select it.
2. In the primary window menu bar, select **Edit—>Properties**.

   or

In the toolbar, click the **Properties** icon.

or

Right-click on the selected device and on the resulting menu select **Properties**.

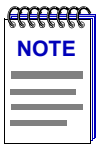

*If you access the Properties window from within a List, Tree, or Map View, you will note that there are two properties-related options available on the menu which appears when you click on the selected device with the **right** mouse button. Select the **Properties** option to launch the appropriate Device Properties window (as displayed above); selecting the **List View, Tree View,** or **Map Properties** option will launch the Properties window for the appropriate view. For more information on tree views, see Chapter 5, **Tree Views**; for more information on maps, see Chapter 6, **Creating Network Maps**.*

Use the appropriate tabbed pages to configure the following display properties. Note that, if you are editing the properties for a device which has already been added to the node database, you may use the **Apply** button to set any changes you make in any tabbed page without closing the window. (The **Apply** button does not appear in the Properties window for newly-added devices.) Clicking **OK** will both save the changes and close the window.

General Properties

The General Properties tab is displayed by default, both when a device is individually added to the database, and when the Device Properties window is launched (see Figure 4-21, page 4-29). Use this window to configure the following properties:

Name	By default, this field will display the MIB-II *sysName* assigned to the device, if one has been assigned; note, however, that editing this name does ***not*** edit the *sysName* value set at the device. The value set in this field will be displayed in the Name column in all list views which contain one, and will also be used to label device icons in Map View windows.
Description	By default, this field will display the MIB-II system description (*sysDescr*) assigned to the device; note, however, that editing this description does not edit the *sysDescr* value set at the device.
Notes	Adding text to the Notes field performs the same function as adding an object note to a device icon or to a list view item (and, in fact, text added to an object note will be displayed here, and vice versa). For more information about adding an object note, see Chapter 6, **Creating Network Maps**.

Device Properties

The Device Properties tabbed page (Figure 4-22) allows you to edit the following properties:

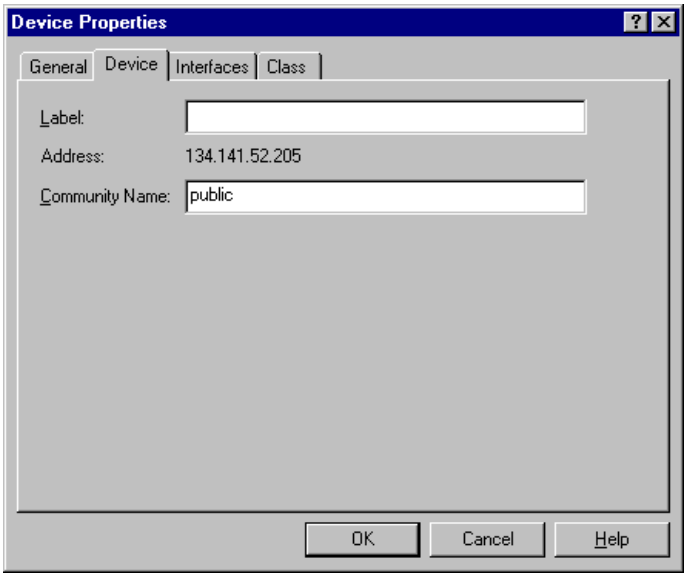

Figure 4-22. Device Properties Page

Label	The value set in this field will be displayed in the Label column in any list view which contains this field. Note that this is *not* the label that will be assigned to a device icon displayed in a Map View. There is no default value.
Address	The IP address assigned to the device. This field is not editable.
Community Name	The community name used to contact the device. This will default to the community name used by the discover script or the community name entered in the Add Device window; you can edit this field if necessary to change your level of access to the device.

Interfaces

The Interfaces tabbed page (Figure 4-23) displays the contact interface on the selected device (i.e., the first interface detected by NetSight Element Manager when the device is discovered or inserted), and both displays and allows you to set its polling parameters. You can also edit certain display properties for that interface.

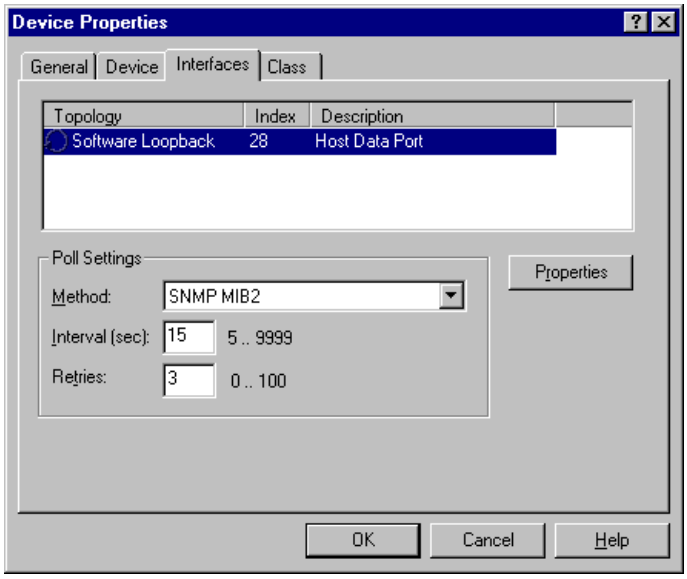

Figure 4-23. Interface Properties

Interface List Box	The list box at the top of the Interfaces page displays the Topology, *IfIndex*, and Description of the contact interface on the selected device.
Poll Settings	The Poll Settings box displays the polling parameters assigned to the contact interface. Even though your managed device may have multiple interfaces, it will only be actively polled on a single interface (the first interface detected when the device is contacted). See **Configuring Node Interface Poll Settings**, page 4-34, for more information.
Properties	Selecting the **Properties** button launches an additional window (see Figure 4-25, page 4-34) that allows you to view display properties for the contact interface, and to configure poll settings for the device. See **Configuring Node Interface Poll Settings**, page 4-34, for more information.

Class Properties

The Class tabbed page (Figure) displays the Model and Enterprise types assigned to the selected device, along with the large (32x32 pixel) and small (16x16 pixel) icons associated with these types. The specific model name of the selected device is also displayed. From this page, you can launch the Node Classes window (Figure 4-27, page 4-37), which allows you to select a different node class; and, via the Node Class Properties window (Figure 4-28, page 4-38), view and change specific parameters assigned to the selected class. You can also use the Node Class Properties window to select the management application that will be launched when the icon is double-clicked

(along with any applicable arguments), and the OID that will be queried to determine the device's type code. For more information on configuring these settings, see **Editing Node Class Properties**, page 4-37.

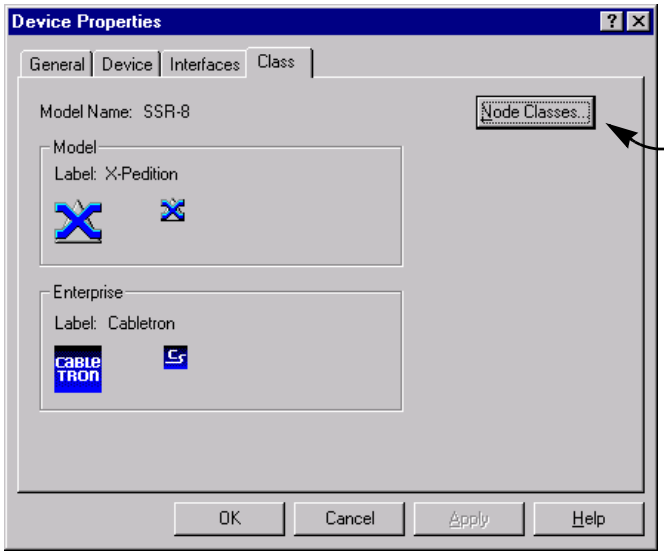

*Click on the **Node Classes** button to display the Node Classes window, which displays the properties associated with each node class, and which allows you to access the Node Class Properties window. See page 4-37 for more information.*

Figure 4-24. Class Properties

Configuring Node Interface Poll Settings

You can configure the parameters which control how NetSight Element Manager polls the contact interface on a selected node via the Poll Settings options in the Interfaces tabbed page in the Properties window (Figure 4-23, page 4-32).

NOTE

*The **Properties** button in the Interfaces tabbed page launches the Interface Properties window, Figure 4-25, with an identical Poll Settings area. Poll Settings information can be configured in either location, with the same effect.*

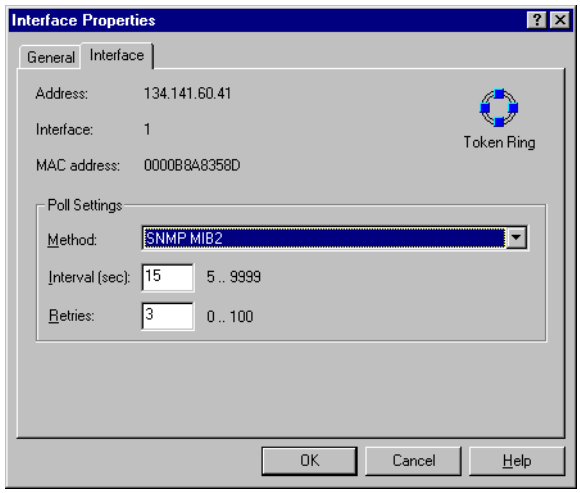

Figure 4-25. The Interface Properties Window

For each node in the database, NetSight Element Manager polls a single interface (for multiple-interface devices, the interface associated with the first valid MAC address detected) via the highest successful polling method (Ping, SNMP MIB I, or SNMP MIB II). By default, the same poll settings are used for all nodes in the database (as determined by the settings in the Polling Options page in the **Tools—>Options** window), and each node's status is updated once during the time it takes to poll the entire database. For example, if you have 500 nodes in your database, and your node polling is set at an interval of 10 seconds with 5 nodes polled per cycle, it will take 1000 seconds (100 poll cycles x 10 seconds), or about 16.5 minutes, to poll the entire database. During that time, the contact status of each polled device will only be updated once; the last known contact status is maintained during the time it takes to refresh the database.

You can alter these database settings via the Interface Properties window, if you wish, but they will take effect **only if the polling interval that you define for the interface is longer than the duration it takes to poll the node database**. Longer poll cycles can be beneficial for non-critical nodes, or nodes for which regular polling would require significant bandwidth cost (for example, those on the other side of a T1 link). By setting a higher polling interval than the node database refresh period, you can reduce the amount of SNMP traffic generated by your NetSight Element Manager station (and any associated bandwidth costs).

If you have a mission-critical node which requires more frequent updates than can be achieved during the database polling cycle, you can keep Chassis Manager running to be constantly appraised of its contact status; the Chassis Manager application uses its own polling mechanism independent of the node database polling in the NetSight Element Manager platform. Refer to the appropriate device-specific ***User's Guide*** for more information on using Chassis Manager; refer to **Adjusting Default Settings** in Chapter 2 for more information on configuring the default Chassis Manager poll settings. You can also issue an immediate status poll at any time by selecting **Manage—>Query Status** from the primary window menu bar, or **Query Status** from the right mouse menu.

To change the poll settings for an individual node:

1. In the **Method** text box, use the down arrow to display a list of selectable polling options: **None**, **Ping**, **SNMP MIB I**, **or SNMP MIB II**. Drag down to select the desired option, and release.
2. In the **Interval** text box, enter the amount of time (in seconds) you want NetSight Element Manager to wait between polls to the interface. The valid range is 5 to 9999 seconds.
3. In the **Retries** field, enter the number of unsuccessful polls that must occur before NetSight Element Manager changes the node status to critical and logs an event. The valid range is 0 to 100.
4. Click **Apply** (if available) to save your changes and leave the Properties window open; or click **OK** to save your changes and close the window.

Three additional values may appear in the **Method** field, although they are not selectable options:

DLM Agent	If the selected device has been configured to be a DLM Agent, the default polled interface will display a Method value of **DLM Agent**, and the polling method will not be editable for that interface.
DLM Client	If the selected device has been configured as a DLM Client and its assigned DLM Agent is functioning correctly, the default polled interface will display a Method value of **DLM Client**, and the polling method will not be editable for that interface.
DLM Client (Direct)	If the selected device has been configured as a DLM Client and its assigned DLM Agent is malfunctioning, the default polled interface will display a method value of **DLM Client (Direct)**, indicating that NetSight Element Manager has taken over polling of the device until the DLM Agent's malfunction is corrected. Again, the polling method will not be editable for that interface.

For more information about DLM and configuring status groups, see the DLM chapter in the ***NetSight Element Manager Tools Guide***.

Editing Poll Settings for Multiple Nodes

You can also edit interface poll settings for multiple nodes via the Multi-Node Properties window (illustrated in Figure 4-26).

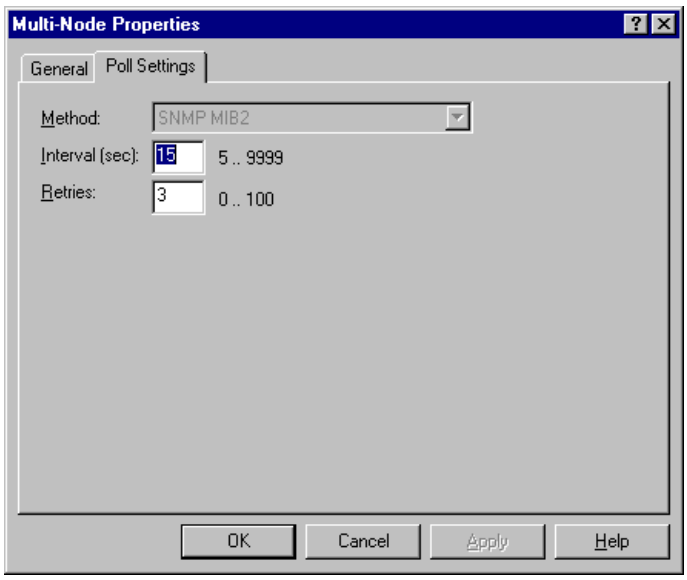

Figure 4-26. The Multi-Node Properties Window

To edit interface poll settings:

1. Open and activate the view which contains the nodes whose interface poll settings you'd like to change, and use *shift-click* or *ctrl-click* to select the appropriate nodes. (Or, if you're using a Map View, you can click and drag to select multiple icons.)
2. In the primary window menu bar, select **Edit—>Properties**.

   or

   In the toolbar, click the **Properties** icon.

   or

   Right-click on any one of the selected nodes, and on the resulting menu, drag down to **Properties**.

Configure the settings as described in the previous section, keeping in mind that only poll settings which are *longer* than the database poll cycle will have an effect on how often the selected nodes' status is updated. Note, too, that you cannot select a polling method via the Multi-Node Properties window; this field will display the currently-set polling method (if all selected nodes are using the same method), or "Mixed" if the selected nodes do not use the same method.

TIP

*You can also add an Object Note to multiple nodes; see **Adding Notes to Multiple Nodes** in Chapter 6, **Creating Network Maps**, for more information.*

Editing Node Class Properties

NetSight Element Manager classifies devices according to what is called a Node Class. In the Node Classes window (Figure 4-27), the properties associated with each node class (model, icon, class, etc.) are displayed in column format. This information can be sorted in forward or reverse order by column, as in your list views, by clicking on the column heading you want the nodes sorted by.

To access the Node Classes window:

1. Open the appropriate device Properties window by selecting the node in any view, and selecting **Edit—>Properties** from the primary window menu bar, clicking the **Properties** icon, or selecting **Properties** from the right-mouse menu.

2. Click the **Class** tab to display the Class tabbed page, then click the **Node Classes** button.

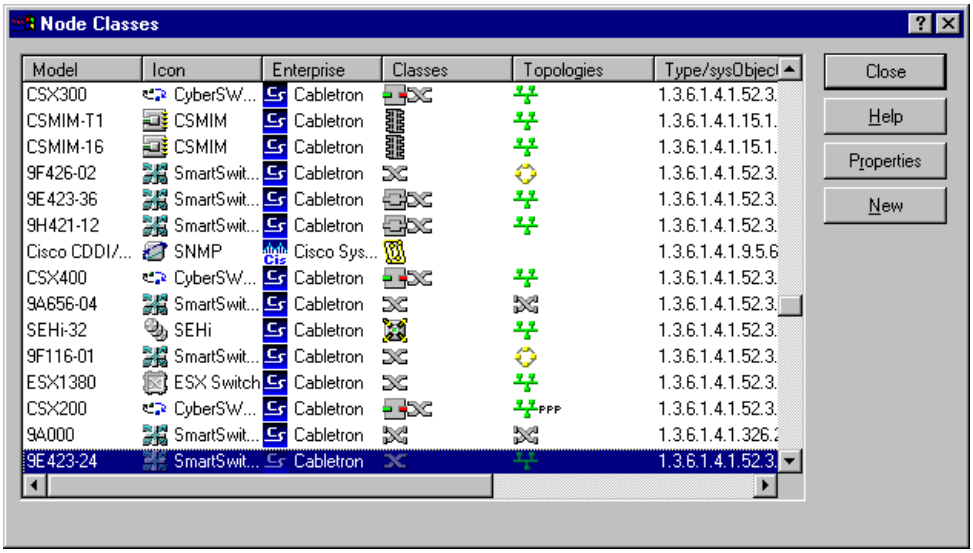

Figure 4-27. The Node Classes Window

The Node Classes window allows you to view the symbols used to represent various nodes and their functional properties; you can edit these default symbol sets or create entirely new ones via the Node Class Properties window.

To edit Node Class Properties for an existing node type:

1. In the Node Classes window, click on the node classes entry you wish to edit, then click on **Properties**.

   or

   Click the **right** mouse button on the node classes entry you wish to modify, and choose **Properties** from the menu that appears. The Node Class Properties window (Figure 4-28) opens.

*Keep in mind that changes made in the Node Class Properties window will apply to **all** nodes of the selected class, not just the specific node whose properties you are modifying. If you wish to launch a specific application for a specific device, you can configure that by customizing the Tools menu (and its accompanying toolbar); see **Customizing the Toolbar** in Chapter 2, **Overview of NetSight Element Manager**.*

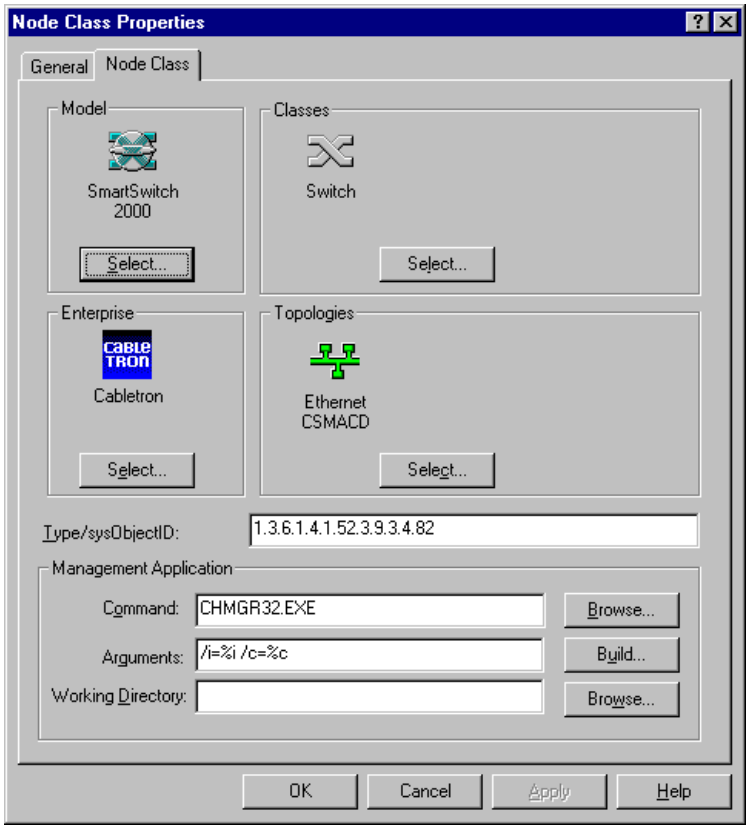

Figure 4-28. The Node Class Properties Window

To create a new node class entry:

1. In the Node Classes window, click on **New**.

   or

   Right-mouse click anywhere in the Node Classes window, and select **New** from the resulting menu.

   The New Node Class window opens. Note that this window is virtually identical to the Node Class Properties window displayed in Figure 4-28, except for the title and the fact that the **General** tabbed page is displayed first by default.

Use the procedures outlined below to edit or create a node class.

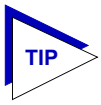

*You can also create new node classes outside of the NetSight Element Manager platform via the **AddDevType** utility; see Appendix A, **Batch Customization via AddTool, AddImage, and AddDevType**, for more information.*

Selecting a New Model, Class, Enterprise, or Topology Symbol

NetSight Element Manager provides a default symbol to represent each of the various model type, class, enterprise, and topology properties that are associated with each node. If you choose, you can change this symbol by selecting a different symbol from the default set, or by creating an entirely new symbol. You can also edit the properties of — that is, change the graphic associated with — any of the symbols provided.

To select a new symbol, click on the **Select** button in the appropriate property area of the window; a Select Symbol window opens, displaying the default set of symbols available for the selected property. Each property — Model, Class, Enterprise, and Topology — supports a different default symbol set, as illustrated in Figure 4-29 (following page).

The process for selecting, editing, and creating symbols is described in detail in **Editing Symbol Properties**, found in Chapter 6, ***Creating Network Maps***.

*Keep in mind that changes made in the Node Class Properties window will apply to **all** nodes of the selected class, not just the specific node whose properties you are modifying If you wish to launch a specific application for a specific device, you can configure that by customizing the Tools menu (and its accompanying toolbar); see **Customizing the Toolbar** in Chapter 2,**Overview of NetSight Element Manager**.*

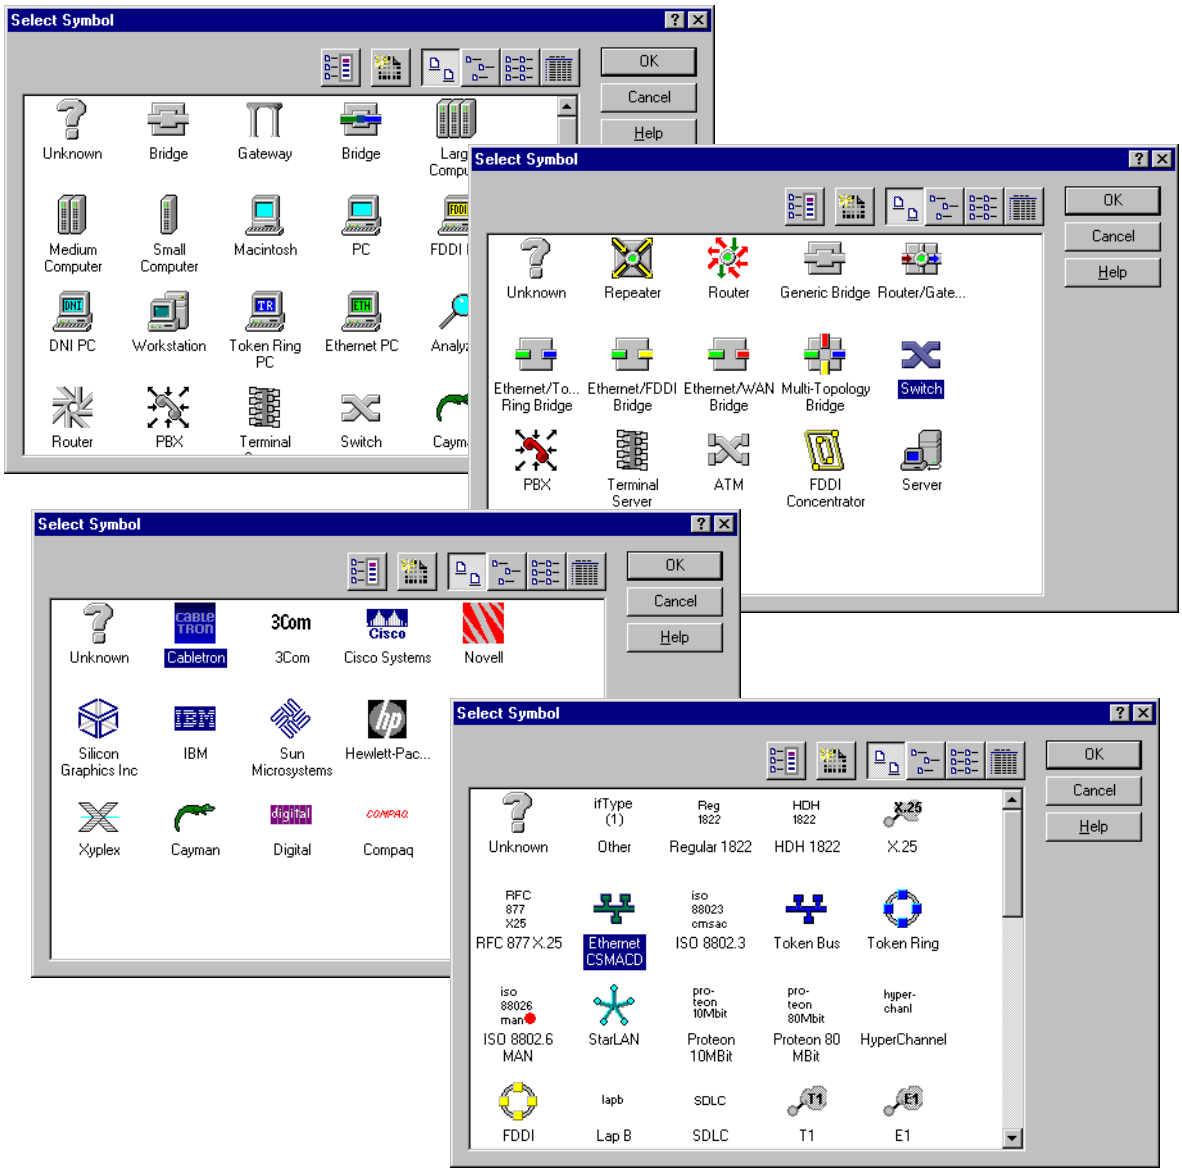

Figure 4-29. Default Symbol Sets

Changing Management Application Information

The Node Class page of the Node Class Properties window also allows you to choose the management application that will be launched for all nodes assigned to a given node class. You specify this via the **Command** field, in which you enter the path and filename of the appropriate executable (*.EXE) file; the **Arguments** field, which allows you to add any necessary or optional arguments to the executable command; and the **Working Directory** field, which allows you to specify (if necessary) the location of any secondary files needed

by the application. These fields can be edited by typing in a new command, argument, or directory, or by using the **Browse** and **Build** buttons, as described in the following sections.

Choosing a Command (Executable) File

The **Command** field in the Management Application area tells NetSight Element Manager which management application is to be launched for a given node class when any one of the **Manage** options is selected. For executables provided by NetSight Element Manager (which all reside in the \NetSight Element Manager x.x\bin directory by default), you need only enter the name of the executable file. For any executable which resides *outside* the default bin directory, you must also enter the full path.

To enter a new command, you can either click in the **Command** text box and type in the name (and path, if necessary) of the executable file you want associated with the node class; or you can use the Browse Executables window, as follows:

1. Click on the **Browse** button next to the **Command** field. The Browse Executables window, Figure 4-30, opens.

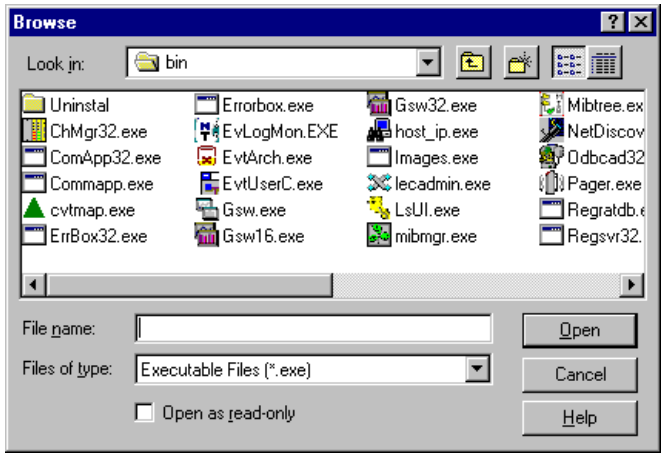

Figure 4-30. The Browse Executables Window

2. The Browse window defaults to the \Netsight Element Manager x.x\bin executables directory; choose an executable file from this list, or change directories as necessary to locate the executable you wish to launch. The executable you have selected will be displayed in the **File name** field at the bottom of the window.

3. When you have made your selection, click **Open**. The Management Application **Command** field will now contain the executable file you have selected.

Building an Argument List

The **Arguments** field allows you to specify any command line parameters (or arguments) that are required by (or permitted by) the selected executable. This field can be manually edited by clicking in it and typing in desired changes, or you can build an argument list using the Build window:

1. In the Management Application area, click on **Build**. The Build window, Figure 4-31, opens.

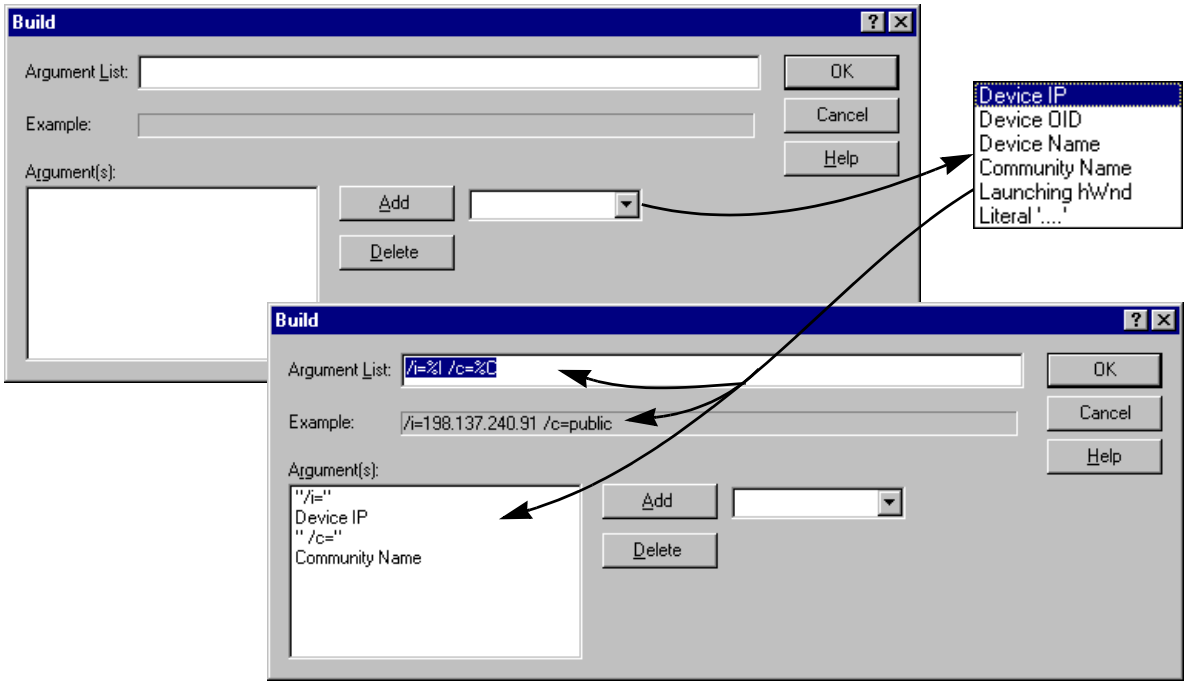

Figure 4-31. The Build Window

2. To add an argument symbol from the pre-defined list provided, click on the down arrow located to the right of the **Add** button, select an argument symbol from the list provided, then click on **Add**. The text description for the argument will be added to the **Arguments** list box, the symbol representing the selected data (%I, for example) will be displayed in the Argument List field, and a sample of the selected argument data (198.137.240.91) will be displayed in the Example field.

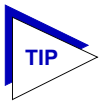

The pre-defined argument list provided lists the argument symbols for data NetSight Element Manager can extract from the central node database for any node against which the selected executable will be launched. Using these symbols allows you to configure a Node Class that can be used for many individual devices. You can also choose to enter specific argument information in the Argument List field (such as a specific community name instead of %C, the symbol representing community name information stored in the database); note, however, that this restricts the use of the Node Class you are configuring to nodes which match the specific argument value you have entered.

3. In the **Argument List** field, add any flags or other command line elements required or permitted by your selected executable. For example, Chassis Manager knows that a **/i** flag indicates that an IP address will follow; **/c** indicates that a community name value will follow. (See below for a list of the flags recognized by some common NetSight Element Manager applications.)
4. Repeat steps 2 and 3 above to add additional arguments as desired. Note that the Build window automatically adds a space (“ “) between each argument you add.
5. If you want any arguments deleted from the field, delete them from the **Argument List** field or from the **Arguments** list box; any deletion in one field will automatically be reflected in the other.
6. When the argument list contains all of the arguments you need, click **OK**. The **Arguments** field of the Node Class page will contain the argument list you have created.

Information NetSight Element Manager can extract about a selected node from the central node database includes:

%I	inserts Device IP from database
%O	inserts Device’s *sysOID* from database
%N	inserts Device Name from database
%C	inserts Community Name from database
%H	inserts the window handle value of the active window

Flags accepted by executables located in \NetSight Element Manager x.x\bin include the following:

chmgr32.exe	Chassis Manager application for device management /i=<ip address> /c=<community string> /p=<.dmf or .mmp file> (for SmartSwitch/MultiSwitch chassis view) /f=<ip address> /c=<community string> (in conjunction with a .dmf or .mmp file, allows users to select from among those files which contain the noted IP address)

In most cases, the flags and related argument values are optional. If you do not specify a necessary value, you will be prompted to supply it when the application launches.

Selecting A Working Directory

If it is required by the executable you have selected, you can also specify a Working Directory. The working directory value sets the location of any secondary files that might be needed by the executable, if those files are not located in the same directory as the executable itself.

To specify a working directory, you can either click in the **Working Directory** text box and type in the path of the appropriate directory; or you can use the **Browse** button, as follows:

1. Click on the **Browse** button next to the **Working Directory** field. A standard Browse for Folder window opens.
2. Select the appropriate directory, and click **OK**. The **Working Directory** field will now contain the directory you have selected.

Changing General Node Class Properties

The other page available in the Node Class Properties window is the General Properties page (Figure 4-32). This page contains fields which allow you to enter or edit a name and description for your node class, as well as any notes you might wish to add. These changes will be applied to all nodes of the selected class.

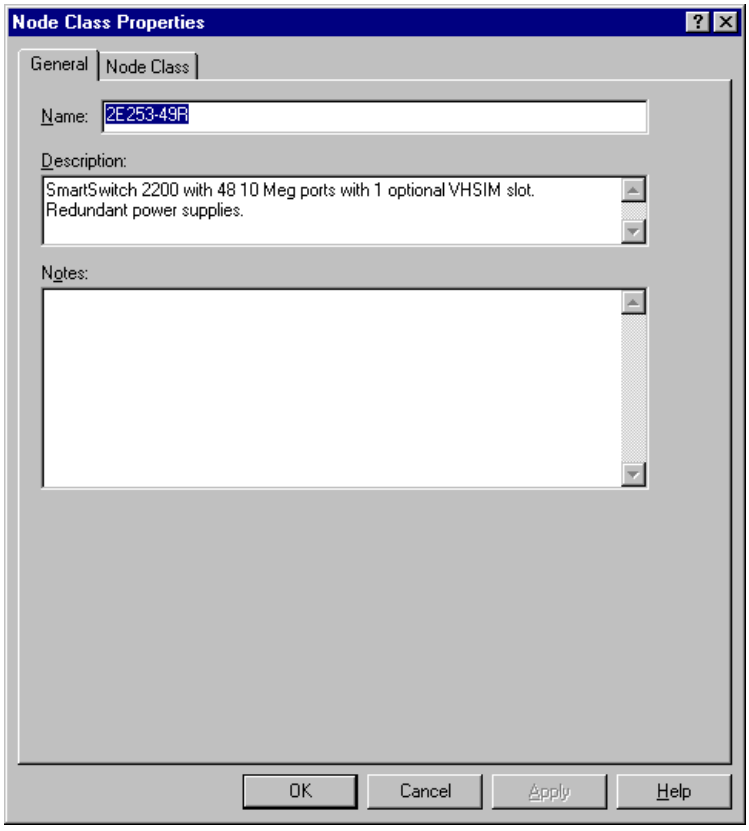

Figure 4-32. The General Properties Window

To enter information for a node class:

1. Click in the **Name** field, and type in a name for the node class. The value entered here is displayed in the Model field of the Node Classes window.
2. In the **Description** field, enter a description for the node class.
3. If there are any notes regarding the node class you wish to keep in the Properties window, type them into the **Notes** field.
4. Click **OK** to save your changes and exit the window.

List Views and Maps

NetSight Element Manager can help you sort out the relationship between your list views and your maps in two ways: by allowing you to select (and, if you wish, delete) any nodes in a view which have *not* been placed on a map; and by helping you locate the map on which a node has been placed.

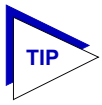

*Both the **Orphan** and **Go to Map** functions require your map files to be open, though they need not be active. NetSight Element Manager will not be able to locate a node's map if that map is closed, and any nodes located on closed maps will be marked as orphans.*

Selecting and Deleting Orphans

Nodes that are displayed in a list view but not placed on a currently open map are recognized by NetSight Element Manager as **orphan** nodes. By using the **Select Orphans** or **Delete Orphans** options from the Edit menu, you can highlight these "orphan" devices in your list view, and, if you choose, delete them.

To select orphan devices:

1. Activate the desired list view and select **Edit —>Select Orphans** from primary window menu bar. Any nodes that are not located on an open map will be highlighted in the list view.

To delete orphans:

1. Activate the desired list view and select **Edit —>Delete Orphans** from the menu bar.
2. The standard Node Delete window opens. Select the appropriate deletion option, and click **OK**.

For more information on the options for deleting and excluding nodes, see **Excluding Nodes from an Individual View**, page 4-18, and **Excluding Nodes from the Central Node Database**, page 4-22. For more information on maps, see Chapter 6, **Creating Network Maps**.

To more accurately select and/or delete orphans from list views, you should have your maps open when using these options. NetSight Element Manager only recognizes a node as being located on a map when the appropriate map is open.

Using the Go To Map Function

You can use the Go to Map function to locate a node on the map on which it has been placed — provided, of course, that the map is open.

1. Open and/or activate the List View which contains the node or nodes whose map you wish to locate, and click to select one node. (The Go to Map options will not be available when multiple nodes are selected.)
2. Select **View—>Go to Map** from the primary window menu bar, or click the right mouse button on the selected node, and select **Go to Map** from the resulting menu.

If the selected node is associated with a map which is open, the correct map window will be activated, and the selected node will be highlighted.

If the selected node is associated with a map which is open but minimized, the minimized map icon will be highlighted.

If the selected node is associated with a map that is closed, or is not associated with any map, a message window opens, indicating that no map containing the selected node could be found.

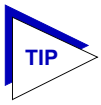

For nodes which are present on multiple maps, only one map will be selected — even if all related maps are open. To locate a node on more than one map, open all maps, use the Go to Map option to locate the node on one map, then close that map and repeat as necessary.

Chapter 5

Tree Views

Accessing a tree view; navigating a tree view with the mouse, keyboard, and Tree View toolbar; tree view properties for nodes and interfaces; the default tree views; creating and editing tree views; excluding nodes

---

Like List Views, Tree Views provide a window through which you can view selected information about the nodes in the central node database. Unlike List Views, however, Tree Views display not only device-level status and configuration, but detailed interface information as well.

About Tree Views

Another way NetSight Element Manager allows you to see, organize, and work with the nodes in your database is through Tree Views. Tree Views, as their name implies, display the contents of your central node database via a tree view control. Each node included in a tree view is displayed at the root of the tree structure, labeled with the appropriate node icon and the node's IP address or name. Clicking on the expand box beside each node icon displays the first level of available information: device-level status, model, and name information. Clicking on any displayed sub-folders (topologies, classes, or interfaces) supplies additional information about the selected node.

As with List Views, you can choose the information about each node that you wish to display and the font used to display it; you can also choose which nodes you wish to display in a view by configuring a filter which selects or omits nodes based on the characteristics you choose. Specific nodes can also be excluded from individual Tree Views, while remaining present in the central node database. Tree Views, like List Views, provide enormous flexibility in the *display* of your central node database without altering its contents.

NetSight Element Manager provides you with several default tree views, including a read-only view called the All Nodes Tree View, which always displays the entire contents of your central node database. Each Tree View can be accessed and (where appropriate) edited via the Tree View Workspace. The workspace also allows you to create new tree views, and organize your views via folders.

The Tree View Workspace

Using the Tree View workspace you can access and edit existing views, create new ones, and organize your views by arranging them in a hierarchy of folders.

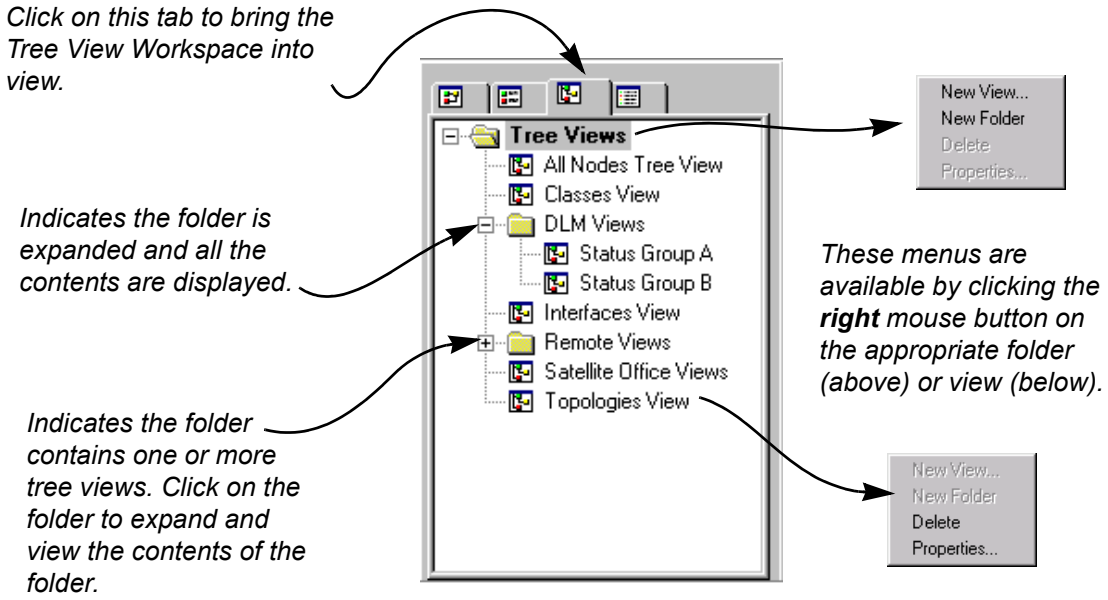
Click on this tab to bring the Tree View Workspace into view.
Indicates the folder is expanded and all the contents are displayed.
Indicates the folder contains one or more tree views. Click on the folder to expand and view the contents of the folder.
Tree Views
All Nodes Tree View
Classes View
DLM Views
Status Group A
Status Group B
Interfaces View
Remote Views
Satellite Office Views
Topologies View
New View...
New Folder
Delete
Properties...
These menus are available by clicking the right mouse button on the appropriate folder (above) or view (below).
New View...
New Folder
Delete
Properties...

Figure 5-1. Getting Around the Tree View Workspace

The default tree views provided with NetSight Element Manager are placed in the root Tree Views folder. All of these default views — except the read-only All Nodes Tree View — can be edited or deleted; you can also add additional views and/or folders.

Adding a Folder to the Workspace

As a means of organizing a collection of Tree Views, you can add a folder to the Tree View page of the workspace; each folder you create will be contained within the root Tree Views folder, and each can contain one or more tree views, as well as additional folders. Note that folders are a workspace display convention only, and they can only be created from within the workspace itself.

To create a folder:

1. Display the Tree View page of the workspace by clicking on the tab in the Workspace View.
2. Click the **right** mouse button on the folder to which you want to add a new folder; that folder will be highlighted (as illustrated in Figure 5-2), and a menu displays.

TIP

To add a folder at the top level of the Tree View hierarchy, click to select the main Tree Views folder.

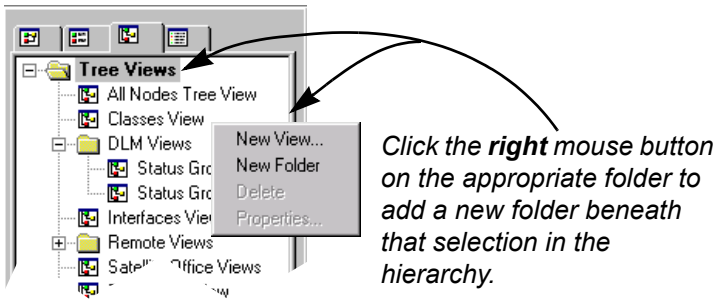

Figure 5-2. Adding a New Folder to the Tree View Workspace

3. In the resulting menu, click to select **New Folder**. A new folder will automatically be added in the branch you have selected.

Renaming a Folder

By default, all new folders are named with the designation **New Folder X**, where X represents a sequential index; you can rename your folders, as follows:

1. Click once to highlight the folder whose name you wish to change, then after a brief pause click again. The folder name will be enclosed in an editable text box (as illustrated in Figure 5-3), and the existing name will be highlighted.

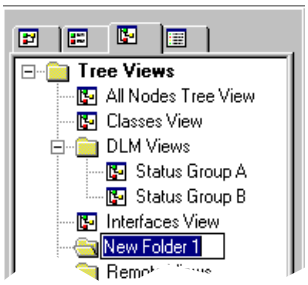

Figure 5-3. Changing a Folder Name

2. Click again to place the cursor anywhere in the name, and edit as desired; or, simply begin typing to replace the highlighted text entirely.
3. Press **Enter** to set your change.

The All Nodes Tree View

The All Nodes Tree View provides the most comprehensive look at the contents of your central node database. This special default view will always appear at the root of your Tree View hierarchy (as displayed in the Tree View workspace), and cannot be edited or deleted; by design, it displays information about every node in the central node database.

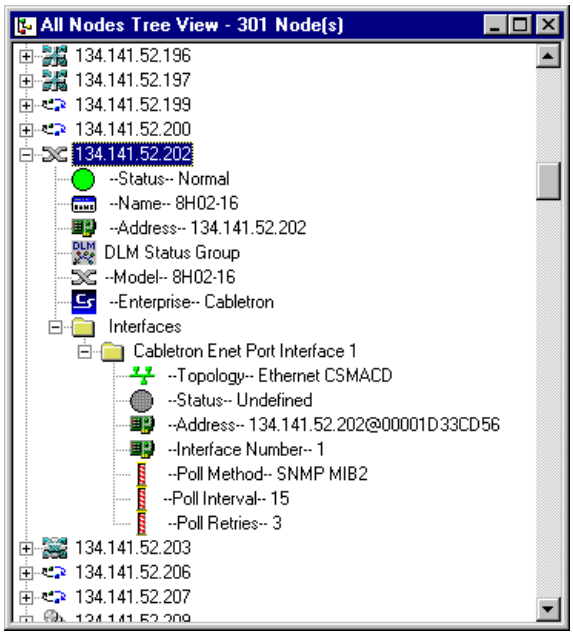

Figure 5-4. The All Nodes Tree View

The All Nodes Tree View display includes seven of the 14 descriptive Node Properties available — including Status, Name, Address, DLM Status Group, Model, Enterprise, and Interfaces — and most of the available Interface information (including Topology, Status, Address, Interface Number, Poll Method, Poll Interval, and Poll Retries). Label, Description, Time of Last Contact, System Up Time, Classes, Topologies, and Notes node properties and Description and Notes interface properties information is omitted (but can be selected for other views). Neither the nodes which are displayed nor the information which is displayed about them can be changed for this view; all of these features, however, are editable in the other default views, and in any new views you create.

TIP

*The one feature you **can** edit in the All Nodes Tree View is the font used to display tree view information. For more information about changing fonts on this and other tree views, see* ***Selecting a Display Font**, page 5-12*.

Other Default Tree Views

As with List Views, NetSight Element Manager provides a set of default views for your convenience. Unlike the List Views, however, these default views don't exclude any *nodes*; they just exclude some of the *information* displayed about the nodes. These default views can be used as is, edited to better suit your needs, or deleted entirely.

Default views provided are:

Classes View

The default Classes view displays basic node information (including Status, Name, Address, Model, and Enterprise), and includes information about the functional Classes supported by the node.

Interfaces View

The Interfaces view displays the same basic set of node information (Status, Name, Address, Model, and Enterprise), and includes detailed information about the Interface through which NetSight Element Manager is communicating with each node.

Topologies View

In addition to the basic node information (Status, Name, Address, Model, and Enterprise), the Topologies view displays the name and icon associated with each topology supported by the selected node.

Navigating a Tree View

There are several means you can use to navigate through a Tree View and display (or hide) the information provided about each node. The basic control mechanism is simple: click on an expand box to display any associated sub-branches of information; click on a collapse box to close the sub-branch display. In addition to using the mouse in the Tree View workspace display, however, you can also use the Tree View toolbar (illustrated in Figure 5-5, page 5-6), the **View** menu, and the arrow keys on your keyboard.

Using the Toolbar and the View Menu

To navigate a tree view via the toolbar or **View** menu, choose the appropriate button or menu option:

Expand One Level	(**View—>Expand**) Expands the selected node or sub-node heading one level, if it has not already been expanded.
Expand Two Levels	Expands the selected node or sub-node heading two levels, if the first and second levels have not already been expanded. If there is only one level left to expand, this button will work the same as the Expand One Level button.

Expand All	(**View—>Expand All**) Expands the entire tree view so that every level of information is displayed for every node.

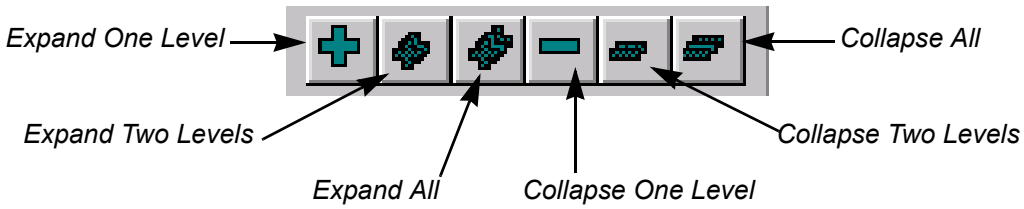

Figure 5-5. The Tree View Toolbar

Collapse One Level	(**View—>Collapse**) Collapses the selected node or sub-node heading one level, if it has not already been collapsed.
Collapse Two Levels	Collapses the selected node or sub-node heading two levels, if the first and second levels have not already been collapsed. If there is only one level left to collapse, this button will work the same as the Collapse One Level button.
Collapse All	(**View—>Collapse All**) Collapses all displayed levels of information so that only the list of nodes (by Name or Address) is visible, along with their model icons.

Using the Keyboard

Once you have displayed and activated a tree view, you can also navigate with the arrow keys and the SPACE BAR:

- The UP and DOWN arrow keys allow you to move one item at a time through a displayed branch of data.
- The RIGHT arrow will expand a collapsed category, and move the cursor to the first item in the expanded branch. (Within the branch, the UP and DOWN arrow keys will select the displayed items.)
- In an expanded display, the LEFT arrow will move the cursor back to the heading of the current branch of information. The LEFT arrow will also collapse an expanded branch.
- The SPACE BAR can be used for changing the highlighted selection within the portion of a Tree View that can be seen in the window (but will not continuously scroll through a Tree View the way UP ARROW and DOWN ARROW will).

Creating and Editing Tree Views

You can create a variety of tree views customized to meet your network management needs, either by editing any of the default views provided (except the All Nodes Tree View) or by creating entirely new views. User-defined filter criteria allow you to select which *nodes* will be displayed; Tree View display properties allow you to select which *information* about those nodes will be displayed.

To create a new tree view:

1. Select **Window—>New Tree View** from the primary window menu bar.

   or

   Click on the toolbar.

   or

   In the Tree View Workspace, click the *right* mouse button on any folder or in any blank area, and select **New View** from the menu that appears.

   The New Tree View window, Figure 5-6, opens.

*The **Node Items** check list allows you to choose which attributes are displayed in the tree view; if you select the Interfaces properties, you can also select which **Interface Items** you wish to display. For more information, see* ***Selecting Node Information to Display**, page 5-8*.

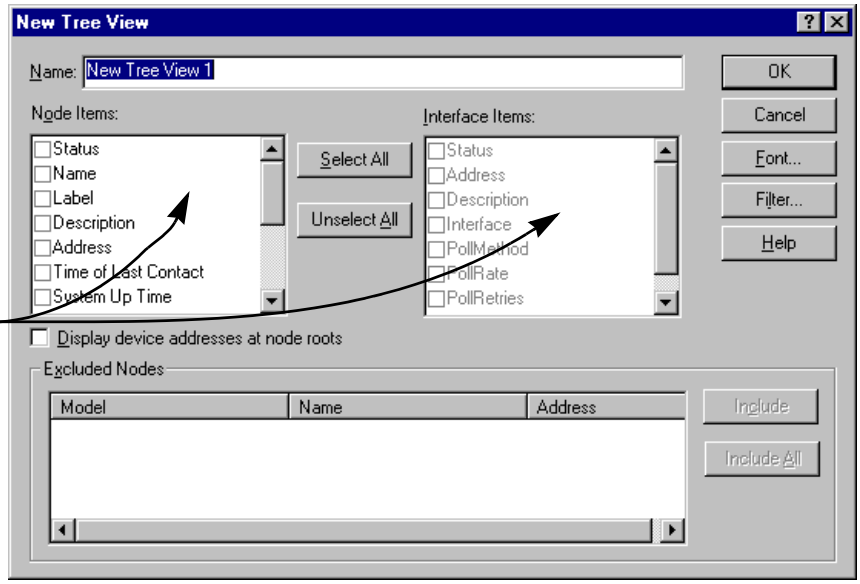

Figure 5-6. The New Tree View Window

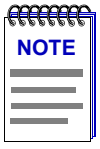

When you add a new tree view by any of the means described above, it will be added to the folder which is currently selected in the Tree View workspace; if you add a tree view by right-clicking in a blank area of the workspace, the new view will be added to the root folder. Once a Tree View has been added, it cannot be moved from folder to folder. All folders and views are listed in the workspace in alphabetical order.

To edit an existing tree view:

1. With the appropriate tree view open and selected in the NetSight Element Manager primary window, select **View —> Properties** from the primary window menu bar.

   or

   In the Tree View Workspace area, click the *right* mouse button on the name or icon associated with the view you wish to edit, and select **Properties** from the resulting menu.

   or

   Click with the *right* mouse button anywhere in the view you wish to edit, and select **Tree View Properties** from the resulting menu.

   The Tree View Properties window will be displayed.

The procedures for creating a new Tree View and editing an existing one are virtually identical; they are described in the following sections.

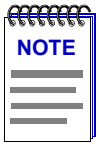

If you open the Tree View Properties window for the All Nodes Tree View, note that all fields are grayed out; remember, the only thing you can change about this view is the font in which its information is displayed. See ***Selecting a Display Font****, page 5-12, for more information.*

Selecting Node Information to Display

The main Tree View Properties window allows you to assign or change the name of your tree view and select the information that will be displayed for each node; you can also choose whether to list nodes by their IP address or host name.

To set up or modify a tree view:

1. By default, a new tree view will be named **New Tree View X**, where X will increment sequentially as necessary to provide a unique name. To change this default (or the original name assigned to any view), highlight the text in the **Name** field and delete, then enter a new name.
2. In the **Node Items** list box, click to select the information you'd like displayed about each node in the view; click again to deselect any option you do not wish to view. Each selection in this list will add a branch of information to your tree view. Each of the available options is described below; you must select at least one property to create or save your view.
3. If you have selected the **Interfaces** node item, the **Interface Items** list box will be activated; use this list box to select the specific interface information you'd like to display.

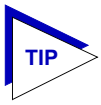

*You can use the **Select All** and **Unselect All** buttons to either select or unselect all available Node and Interface items.*

4. By default, a tree view will list nodes by their assigned Name; if you prefer devices to be listed by IP address, click to select the **Display device addresses at node roots** option.

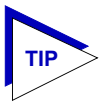

Unlike List Views, Tree Views cannot be sorted; nodes will be displayed in numeric order (if IP addresses are displayed at the node root), or in alphabetical order (if nodes are displayed by Name). Changing the root display (from IP address to Name, or vice versa) will change the order of nodes in the view. By default, a node is assigned its sysName *(if one has been configured) for a Name value; you can also choose to assign a different Name value via the node's Properties window. See Chapter 4, **List Views**, for more information about configuring node properties.*

A tree view configured in this way will display the information you have selected for *all* nodes in the central node database, using the *default* display font. To change the node types that will be included in your view, use the **Filters** option (described in **Selecting the Nodes to Display: Using Filters**, page 5-13); to change the display font, use the **Font** option (described in **Selecting a Display Font**, page 5-12).

*For more information about the view-specific Excluded Nodes list displayed at the bottom of the Tree View Properties window, see **Excluding Nodes from an Individual View** in Chapter 4, **List Views**.*

Node Properties Options

Each tree view can include up to 14 branches of information about the nodes it displays, some of which can be further expanded. Information available for display includes:

Status

Indicates the status of NetSight Element Manager's contact with the node. The possible states are:

- **Critical** (Red) - The node is in critical condition; it is no longer in contact with NetSight Element Manager.
- **Warning** (Yellow) - NetSight Element Manager has received a trap from the node. The icon will remain yellow until the trap is acknowledged.

If any node in your central node database is in a Critical or Warning state, this status condition will be propagated right up to your desktop via the minimized NetSight Element Manager icon — the Task Bar icon representing NetSight Element Manager will change from the inactive window title bar color to the active window title bar color.

- **Unknown** (Blue) - NetSight Element Manager has not yet contacted the node, or has not determined whether that contact was successful. A node icon is blue only temporarily, while NetSight Element Manager determines whether contact can be made: if contact is successfully made, the icon will turn green; if not, it will turn red. Performing an individual node status query (using the **Manage—>Query Status** option from the primary window menu bar, or the **Query Status** option available from the right-mouse menu) will temporarily turn a node status icon blue, while the status query is being made.
- **Normal** (Green) - NetSight Element Manager is successfully communicating with the represented node, and the node is operating properly on the network.
- **Undefined** (Gray) - Polling has been disabled, and the node's current status is not known.

Name

This column displays the user-defined name of the node. You can assign a name for the node in the **General** properties page in the Device Properties window. By default, devices will display any configured MIB II *sysName* value. You can choose to have the name displayed at the root of the tree, see **Selecting Node Information to Display**, page 5-8, for more information.

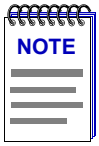

*Changing a node's name via its Properties window changes **only** the information in the central node database; no name changes are set at the device. If you wish to leave the MIB II* sysName *displayed, **and** assign an additional label, use the Label field described below. For more information on configuring node properties — including setting both Name and Label values — see **Editing Device Properties** in Chapter 4, **List Views**.*

Label

Displays a user-set alias for a node; no default label is provided.

Description

Displays the descriptive information assigned to the node. By default, this field will display the *sysDescr* value configured at the device (which typically includes the model name of the node and the version number of its firmware image), but you can assign a new description to a node via its Properties window. For more information about configuring node properties, see Chapter 4, **List Views**.

Address

Displays the IP (Internet Protocol) address for the device.

Time of Last Contact

Lists the date and time of NetSight Element Manager's last successful contact with the node.

System Up Time

The length of time the device has been operating since its last boot up, displayed in a days, HH:MM:SS format.

DLM Status Group

If the device is part of a Distributed LAN Management status group, it will be indicated here. Both the name assigned to the DLM group and the node's position in that group ("DLM Agent" or "DLM Client of (IP Address)") will be displayed. For more information on configuring and using Distributed LAN Management functions, see the **DLM Status Groups** chapter in the ***Tools Guide***.

Model

Indicates the model name of the device, along with the icon which represents that model type.

Classes

Displays the functional classes supported by the selected node (Router, Bridge, Switch, ATM, etc.), represented by the appropriate class icon(s).

Enterprise

Displays the enterprise or vendor that manufactured the selected node.

Topologies

Displays the topology (or topologies) supported by the node.

Interfaces

Displays the interfaces branch of information for the selected node, which includes a variety of information about the interface through which NetSight Element Manager is communicating with the selected node. The specific interface information that can be selected for display (via the Interface Items list box) is described in the following section.

Notes

Displays the text of any notes you have assigned to the node. For more information on Notes, see Chapter 6, **Creating Network Maps**.

Interface Properties Options

If you have selected the Interfaces node property for display, you can choose to display a variety of data specific to the interface through which NetSight Element Manager is communicating with each displayed node.

Simply choosing the Interface Node Item displays a description of the contact interface, which serves as the label for the folder which contains the selectable Interface Items. Also displayed automatically is the first item in the interfaces branch contained in the folder:

Topology	Displays the network topology supported by the contact interface, along with an icon representing that topology.

Additional interface properties available for selection are:

Status	Indicates the status of NetSight Element Manager's contact with the node. Status conditions and their corresponding colors are described beginning on page 5-9.

Address	Displays the IP address assigned to the node, and the MAC address of the contact interface.
Interface Number	Display the *ifIndex* value assigned to the contact interface.
Poll Method	Indicates the type of polling currently being used to contact the node.
Poll Interval	The configured interval between interface polls, in seconds.
Poll Retries	The number of polls that must fail before NetSight Element Manager declares communication with the interface has failed.
Notes	Displays the text of any notes you have assigned to the interface. These notes can also be edited via the General tab in the Interface Properties window. For more information on Notes, see Chapter 6, **Creating Network Maps**; for more information on the node properties window, see Chapter 4, **List Views**.

Selecting a Display Font

The Font Window allows you to change the appearance of your tree view by allowing you to select the font style and size for any displayed text. You can change the font setting for *any* tree view, including the All Nodes Tree View.

1. Click on **Font** from the appropriate New Tree View or Tree View Properties window. The Font window, Figure 5-7, will be displayed.

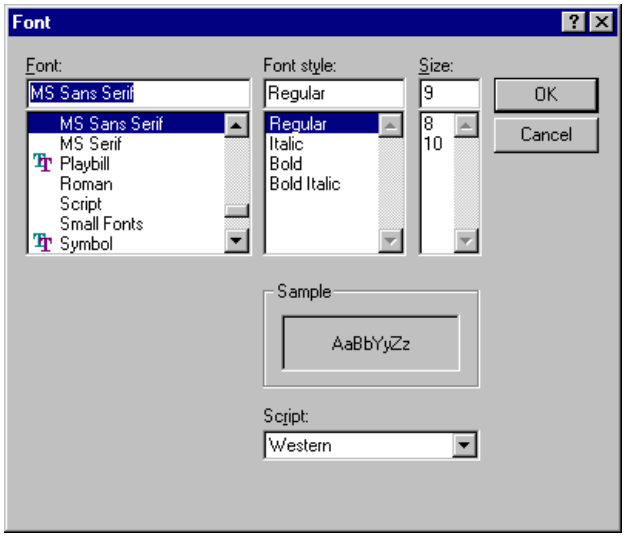

Figure 5-7. Font Window

Change your font settings as desired using the **Fonts**, **Font Style**, and **Size** fields, then click on **OK** to set your changes.

Selecting the Nodes to Display: Using Filters

Using the View Filter window, you can set up filters that determine which of the nodes stored in the central node database will be displayed in your tree view. You can filter nodes based on specific **Subnet**, **Status Group**, **Enterprise**, **Class**, and/or **Topology** attributes.

To open the View Filter window:

1. Click on **Filter** from the New Tree View or Tree View Properties window. The View Filter window, Figure 5-8, will be displayed.

The check boxes above the filter categories allow you to either select or de-select all of the specific attributes listed. You must leave at least one attribute selected in each category; to logically eliminate a category from the filter, select all of its attributes.

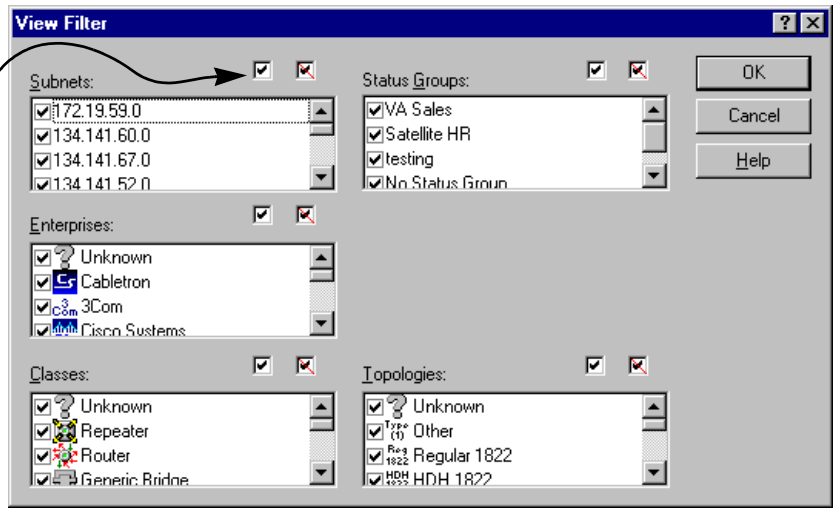

Figure 5-8. View Filter Window

Filter Categories

Available filter categories are:

Subnets

NetSight Element Manager automatically creates a subnet entry each time a node is added to the database; subnet entries can also be defined via the Discover application and/or via the **View—>Subnets** option in the primary window menu bar. All subnets which have been created will be displayed in this list box; use this list to select the subnets whose nodes you wish to display in your tree view.

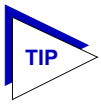

*See Chapter 3, **Discovering Nodes**, for more information on creating and using subnet entries.*

Status Groups

Any DLM Status Groups you have created will be displayed in this list box. You can choose to display only those nodes which are part of one or more selected status groups, or those which are part of any defined status group.

Enterprises

The Enterprise category allows you to filter nodes by hardware or software manufacturer. The vendor codes available for filtering are Cabletron, 3Com, Cisco Systems, Novell, Silicon Graphics, IBM, Sun Microsystems, Hewlett-Packard, Microsoft, Xyplex, Cayman, Digital, Compaq, and Enterasys. You can also choose to filter on nodes whose vendor code is unknown.

Classes

Lists the functional classes that NetSight Element Manager can detect. Classes available for filtering include:

- Repeater
- Router
- Generic Bridge
- Router/Gateway
- Ethernet/Token Ring Bridge
- Switch
- Ethernet/FDDI Bridge
- Ethernet/WAN Bridge
- PBX
- Multi-Topology Bridge
- FDDI Concentrator
- Terminal Server
- Server
- ATM

Nodes which support more than one functional class will be displayed in any view for which one of its classes is selected. You can also choose to filter on nodes whose functional class is unknown.

Topologies

Lists the topologies that NetSight Element Manager will recognize. Topologies listed include:

- Regular 1822
- HDH 1822
- Local Talk
- RFC 877 X.25
- Ethernet CSMACD
- ISO 8802.3
- ISO 8802.6 MAN
- ISO 88022 LLC
- 3 Mbit Ethernet
- StarLAN
- Token Bus
- Token Ring
- HyperChannel
- FDDI
- Lap B
- T1
- T-3
- E1
- Basic ISDN
- Primary ISDN
- PPP
- Serial Line IP
- CLNP/IP
- XNS/IP

- Software Loopback
- Ultra Tech
- SMDS
- Proteon 80Mbit
- Proteon 10Mbit
- ATM
- AAL5 over ATM
- RS-232
- Parallel Port
- Arcnet
- Arcnet Plus
- SMDS DXI
- SONET
- SONET Path
- SONET VT
- SDLC
- V.35
- Frame Relay
- Frame Relay Service
- HIPPI
- HSSI
- X.25
- Proprietary Multiplexing
- X25PLE
- M10X25
- Proprietary Virtual/Internal
- Generic Modem
- Proprietary PPP Serial
- SMDS Intercarrier Interface
- Other

As with classes, nodes which support more than one topology will be displayed in any view for which one of its topologies is selected. You can also choose to filter on nodes whose topology is unknown.

Creating the Filter

When you create a filter, you select the specific attributes in each category that define the nodes you *want* to display, and de-select the attributes that define nodes you do *not* want to display. Filter categories are linked by logical ANDs, and attributes within categories are linked by logical ORs; this means that, in order to be displayed in a filtered view, a node must match *at least one* selected attribute in *each* category. You must leave at least one attribute selected in each category.

To specify a view filter:

1. In each filter category list box, click on the check box next to the specific attribute to include or exclude associated nodes from the tree view.

You can also use the select and deselect check boxes located above each filter category list box to either select or deselect all of its specific attributes; remember, however, that you must leave at least one attribute selected in each category in order to create your filter.

2. Click on **OK** to set the filter parameters.

Once you have finished configuring all of the parameters for your tree view, click **OK** in the New Tree View window or Tree View Properties window to create the new tree view or modify an existing view.

Excluded Nodes

NetSight Element Manager's architecture — which stores all node information in a central database, and lets you choose how to view it — provides you with great flexibility in the management and storage of node information. One of the ways you can manage your collected node information is through the use of Excluded Nodes. There are two levels of node exclusion: exclusion from an individual view, and exclusion from all views. Exclusion from an individual view removes a node from that view and places it on the Excluded Nodes list displayed in the Tree View Properties window; exclusion from all views removes the selected node from the central node database into the Excluded Nodes list.

For more detailed information on the process and effect of node exclusion, see **Excluded Nodes** in Chapter 4, **List Views**.

Tree Views and Maps

NetSight Element Manager can help you sort out the relationship between your tree views and your maps by helping you locate the map on which a node has been placed.

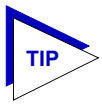

The Go to Map function requires your map files to be open, though they need not be active. NetSight Element Manager will not be able to locate a node's map if that map is closed.

Using the Go To Map Function

You can use the Go to Map function to locate a node on the map on which it has been placed — provided, of course, that the map is open.

1. Open and/or activate the Tree View which contains the node whose map you wish to locate, and click to select one node.
2. Select **View—>Go to Map** from the primary window menu bar, or click the right mouse button on the selected node, and select **Go to Map** from the resulting menu.

If the selected node is associated with a map which is open, the correct map window will be activated, and the selected node will be highlighted.

If the selected node is associated with a map which is open but minimized, the minimized map icon will be highlighted.

If the selected node is associated with a map that is closed, or is not associated with any map, a message window opens, indicating that no map containing the selected node could be found.

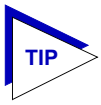

For nodes which are present on multiple maps, only one map will be selected — even if all related maps are open. To locate a node on more than one map, open all maps, use the Go to Map option to locate the node on one map, then close that map and repeat as necessary.

Chapter 6

Creating Network Maps

Building a map hierarchy; about the Map View; creating new maps; adding submaps and GoTo symbols; generating submaps automatically; adding nodes to a map; importing MRM maps; customizing your maps; aligning map objects; saving maps

---

NetSight Element Manager's mapping capabilities provide a powerful means for creating logical and/or physical representations of your network. Unlike List and Tree Views, Map Views are not populated automatically; you can create as many maps as you need, and add and arrange node icons in any way you like. Maps can also be nested via a submap hierarchy, which can be displayed at a glance via the Map View page in the workspace. You can add device icons directly to your maps, or simply drag them in from any existing view; you can also add a variety of symbols to your map, and take advantage of Microsoft's Object Linking and Embedding (OLE) technology to add objects created by other applications.

Working with Maps

NetSight Element Manager creates and maintains maps without assigning any default hierarchical structure to them: each map is created and maintained as an independent file. If you wish, however, you can define relationships among any number of independent maps by establishing links between them. These links set up a sequence of status reporting that results in a *de facto* hierarchical structure defined by the direction in which maps report their status to other linked maps.

NetSight Element Manager allows you to create a nested hierarchy of maps via the use of **submap icons**. In this nested hierarchy — which can be adjusted and changed as needed — each map reports its status (and the status of its subparts) to all related maps.

About the Map View

The Map View portion of the NetSight Element Manager workspace (illustrated in Figure 6-1) provides a bird's-eye view of the submap hierarchy you have created among your maps via a standard indented tree structure: maps can be grouped in **Folders**; folders can contain **Maps** and/or additional folders; and each map can contain one or more **Submaps**.

The relationships between maps and folders are defined when the maps and folders are created, and they cannot easily be changed. For example, if you want to group certain maps within a folder, you must first create the folder, then be sure to create the maps from within that folder.

The relationships between maps and submaps are more flexible: you can link two existing maps via a submap icon, or create a new submap when you add a submap icon; you can also change map/submap relationships simply by moving (or removing) a submap icon. Creating submap relationships will move maps from folder to folder, when necessary; breaking those relationships places the former submaps at the root of the workspace hierarchy. All of these changes will be reflected in the Map View workspace.

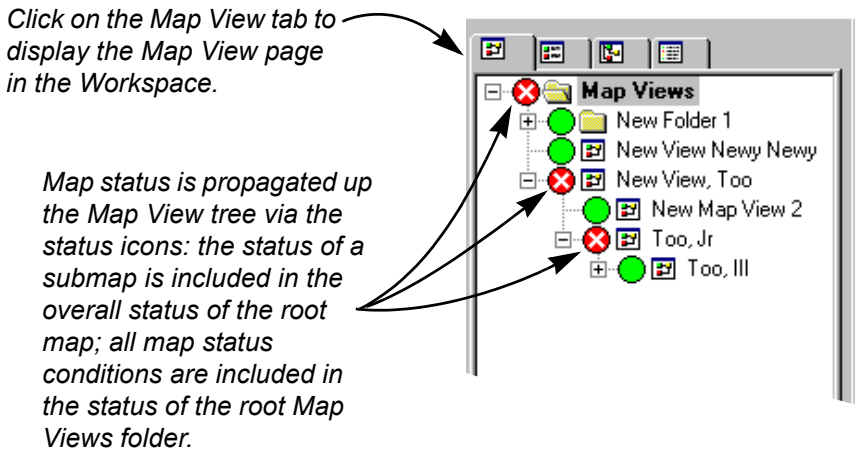
Click on the Map View tab to display the Map View page in the Workspace.
Map status is propagated up the Map View tree via the status icons: the status of a submap is included in the overall status of the root map; all map status conditions are included in the status of the root Map Views folder.
Map Views
New Folder 1
New View Newy Newy
New View, Too
New Map View 2
Too, Jr
Too, III

Figure 6-1. The Map View Page

Each element in the nested submap hierarchy reports its overall status to the next level up in the hierarchy, as illustrated in Figure 6-1 above. The map named **Too, Jr** has at least one element in a critical state, as indicated by the red status icon; that red status is propagated up to the map named **New View, Too**, of which Too, Jr is a submap. That red status is further propagated up to the main **Map View** folder, which is always the root folder for all maps and submaps. Map status is displayed and propagated in the workspace whether your maps are open or closed. Note that the plus and minus symbols indicate whether the selected branch is fully displayed or whether additional submaps and/or folders are currently hidden from view. To fully display a branch, click on the plus symbol; to close it, click on the minus symbol.

From within the Map View workspace, you can create, delete, or rename folders and maps, and delete or rename submaps; you can also launch the Properties window for any map or submap. Each of these functions is described in the following sections.

About Submap Icons

Using Submap icons to link maps has two results: first, the map defined as the submap reports its status — and that of its component parts — to its defined root map; and second, the hierarchical arrangement that results from defined submap relationships is graphically displayed in the Map View workspace. For example, if you have created one map called BUILDING_A, and a second map called FLOOR_1, you can link those two maps into a hierarchical structure by adding a submap icon representing FLOOR_1 to BUILDING_A; that submap icon will change colors to reflect the current status of FLOOR_1, and that submap relationship will be displayed in the tree structure shown in the Map View. You can then either widen the hierarchy by creating additional maps that report their status directly back to BUILDING_A, or deepen the hierarchy by creating additional maps that report their status to FLOOR_1 (which in turn reports its collective status back to BUILDING_A). Again, all of these relationships are displayed in the Map View tree structure.

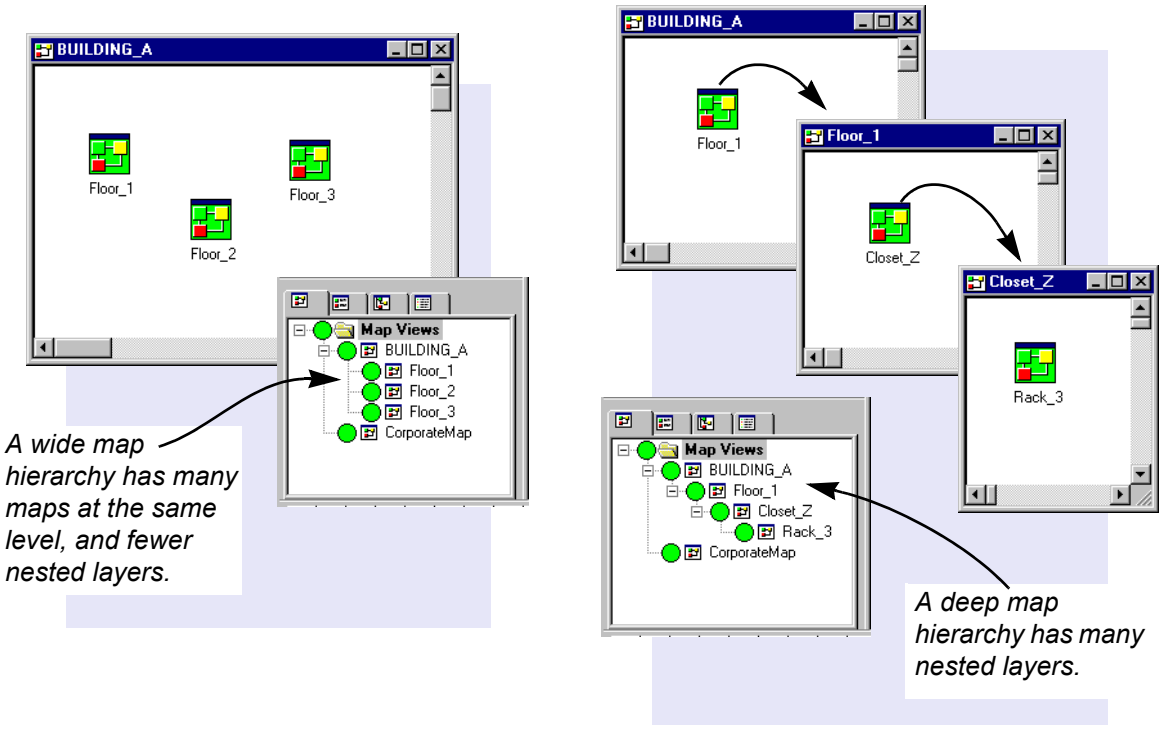

Figure 6-2. Building a Map Hierarchy

There is only one restriction to submap relationships: a map cannot be defined as a submap of more than one root map. In the above example, then, FLOOR_1 cannot be made a submap of any other root map while it is defined as a submap of BUILDING_A. If you try to add a second submap icon representing the same (FLOOR_1) map, the submap relationship will change — for example, adding FLOOR_1 as a submap of BUILDING_B removes its submap relationship with BUILDING_A. This change will occur without warning, but is also reflected in the workspace.

When creating submap relationships, place the appropriate submap icon in the appropriate map. You can use a submap icon to link two existing maps, or create a new map when you place the submap icon. Similarly, changing an established hierarchy can be accomplished by moving or removing the appropriate submap icons. You can even create a new map and place it at the top of your hierarchy just by adding the appropriate submap icons.

You can also create submaps based on subnets automatically; see ***Generating Submaps Automatically****, page 6-15, for details.*

About GoTo Symbols

Using GoTo symbols to link maps allows you to create a more free-form relationship among maps — one that can also be linked and unlinked at will. Placing a GoTo symbol into a map provides a quick and easy means of accessing the map represented by the symbol, and allows you to link maps outside of the status reporting hierarchy. GoTo symbols report the status of the maps they represent, but this status is not included in the overall status of the map in which the GoTo symbol has been placed. The relationships defined by GoTo symbols are not displayed in the Map View tree structure; there is no limit to the number of GoTo symbols you can add representing a single map.

Map and Device Status Conditions

Each entity in a nested hierarchy (submap, map, or folder) reports its overall status (and the status of its subparts) to the next layer up in the tree structure, as indicated by the color of the status icon located to the left of each map or folder entry in the Map View.

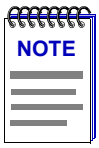

GoTo symbols also report the overall status of the maps they represent, but the status reported by a GoTo symbol is not part of the overall status of the map in which it is placed, and will not be propagated up the tree hierarchy.

Status conditions as reported by each map — listed here in decreasing order of severity — are defined as follows:

Critical (Red)	At least one element on the represented map is reporting a critical condition; a node in critical condition is no longer in contact with NetSight Element Manager.
Warning (Yellow)	At least one element in the represented map is reporting a warning condition. A node icon turns yellow to indicate that NetSight Element Manager has received a trap from the represented node, and remains yellow until the trap is acknowledged.

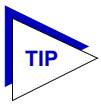

If any node in your central node database is in a Critical or Warning state, this status condition will be propagated right up to your desktop via the minimized NetSight Element Manager icon — the Task Bar icon representing NetSight Element Manager will change from the inactive window title bar color to the active window title bar color.

Unknown (Blue)	NetSight Element Manager has not yet been able to contact at least one element in the associated map, and has not yet determined whether or not that contact will be successful. A node or map icon is blue only temporarily, while NetSight Element Manager determines whether or not contact can be made: if contact is successfully made, the icon will turn green; if not, it will turn red. Performing an individual node status query (using the **Manage—>Query Status** option from the primary window menu bar, or the **Query Status** option available from the right-mouse menu) temporarily turns a node icon blue, while the status query is being made.
Normal (Green)	No elements in the map are reporting warning or critical conditions; a map with no elements will also report this status. NetSight Element Manager is successfully communicating with the represented node, and the node is operating properly on the network. Green is also the default color for an undefined icon (represented by a question mark): one with which no IP address is currently associated.

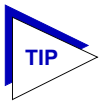

*If you exclude or delete a node from the central node database, its icon will automatically be removed from any open maps; if it appears on any closed maps, its node-specific icon will turn into an undefined (question mark) icon, but remain in place. Note that re-including nodes from the Excluded Nodes database restores them to all appropriate List and Tree Views, but does **not** restore them to maps.*

Undefined (Gray)	Polling has been disabled. This gray status does not propagate to the map.

Building Your Network Map

Before you begin creating maps, it's useful to spend a few moments thinking about the organization of your network and how best to create a visual representation of that organization for the purposes of network management. If you wish to organize your maps in folders, it's important to know in advance which folders to create, since you must add new maps from within a selected folder to create that relationship (and maps cannot easily be moved from one folder to another). You may also want to decide whether to arrange your networking devices by their physical location, by their logical grouping, or by some combination of the two.

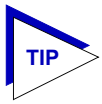

Each map file is saved in two versions: a .MAP file and a .MRM file. The .MAP file is the version read by NetSight Element Manager. The .MRM file is a text version that can be re-imported into NetSight Element Manager if the .MAP file is corrupted or lost. You can move a map from one folder to another by taking advantage of these duplicate files: simply delete the .MAP file, click to select the folder you want the map stored in (in the workspace), and import the associated .MRM file. Deleting .MAP files from the workspace does not remove the associated .MRM file.

Map-folder relationships will also change as necessary to reflect submap relationships: if you make a map from one folder a submap of a map in another folder, the submap will move from its original folder to the one containing its root map. If you break that submap relationship, however, the former submap moves to the root of the tree hierarchy — not back into its original folder.

In addition to a nested hierarchy of maps and submaps, you may also want to add a few GoTo symbols; note, however, that while these symbols display the status of the maps they represent, this status is not propagated to the map which contains the GoTo symbol.

You can create new folders only from within the Map View workspace; you can add new submaps and GoTo symbols from the **Insert** menu or the toolbar; and you can add new maps from the **File** menu or from within the Map View workspace. You can also generate submaps automatically. Each of these processes is described in the following sections.

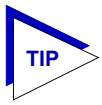

The instructions provided in this chapter frequently reference menu and toolbar options which may not be available if no maps are open. If you don't find a menu option where you expect it to be, or if the toolbar option you want to use is not active, check to make sure the Map View page of the workspace is displayed, and that at least one map file is open and active.

Adding a Folder

As a means of organizing a series of maps, you can add a folder to the Map View page of the workspace; each folder you create will be contained within the base Map View folder, and each can contain one or more maps, submaps, and additional folders. Each will also

display the collective status of the maps it contains. Note that folders are a Map View workspace display convention only, and they can only be created from within the Map View workspace itself.

1. Display the Map View page of the workspace by clicking on the tab in the Workspace View.
2. Click the **right** mouse button on the folder to which you want to add a new folder; that folder will be highlighted (as illustrated in Figure 6-3), and a menu opens.

TIP

To add a folder at the top level of the Map View hierarchy, click to select the main Map Views folder.

3. In the resulting menu, click to select **New Folder**. A new folder will automatically be added in the branch you have selected.

Renaming a Folder

By default, all new folders are named with the designation New Folder X, where X represents a sequential index; you can rename your folders, as follows:

1. Highlight the folder whose name you wish to change, then after a brief pause click again. The folder name will be enclosed in an editable text box (as illustrated in Figure 6-3), and the existing name will be highlighted.

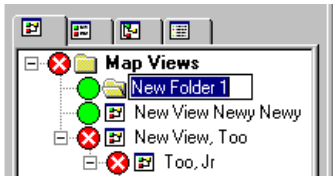

Figure 6-3. Changing a Folder Name

2. Place the cursor anywhere in the name, and edit as desired; or, simply begin typing to replace the highlighted text entirely.
3. Press **Enter** to set your change.

Adding a New Map

You can create a new map file via the primary window **File—>New Map** menu option, via the New Map toolbar button , or via the menu available by clicking the **right** mouse button within the Map View workspace. To add a map file to an existing folder,

you must click the **right** mouse button on the folder of choice and select **New View** from the resulting menu (as described below), or make sure the folder you want to add the map to is highlighted in the Map View before using the File menu or toolbar options.

TIP

You can also create a new map when you add a Submap or GoTo icon; see ***Adding a Submap****, page 6-13, or* ***Adding a GoTo Symbol****, page 6-18.*

To create a new map:

1. In the Map View window, select the folder into which you wish to add the new map. If you want to add the new map to the top level of the Map View hierarchy, select the main Map Views folder, or highlight any existing map.
2. In the primary window menu bar, select **File—>New Map**.

   or

   In the toolbar, click the **New Map** icon.

   or

   In the Map View workspace, click the **right** mouse button on the selected folder, and select **New View** from the resulting menu.

The New Map View window, Figure 6-4, opens.

Figure 6-4. The New Map View Window

Once you have created the new map, it displays in the appropriate place in the Map View tree hierarchy.

The New Map View window allows you to set a variety of display and print options for your map. To create a map file, you need only assign a unique name; the name and any other map parameters can be changed later by accessing the Map Properties window (see **Editing Map Properties**, page 6-12, for more information). The following sections describe how to set map parameters.

Setting Basic Map Options

NetSight Element Manager allows you to set a variety of display and print options for each map or submap you create. You can set these options when you first create a map or submap (via the New Map View window illustrated in Figure 6-4, above) or at any time after a map or submap has been created (via the Map View Properties window, described in **Editing Map Properties**, page 6-12). Among the properties that you can set are its background image and/or color, background grid settings, and print options. You can also select a font for icon labels, and lock the map so that it cannot be altered or deleted.

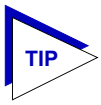

*For more information on the Lock feature, see **Protecting Your Maps from Editing: The Lock Feature**, page 6-60.*

When you have finished setting map parameters, click on **OK** to set your changes, or on **Cancel** to cancel them and exit the New Map View window.

Selecting a Background Color or Image

You can assign a color or a bitmap image to the map's background, as follows:

1. To choose background color for the Map window:
    a. Select the **Color** option; note that when you select the **Color** option, the **Bitmap** fields are grayed out.
    b. Click **Set Color**. The Color window (Figure 6-5) opens.

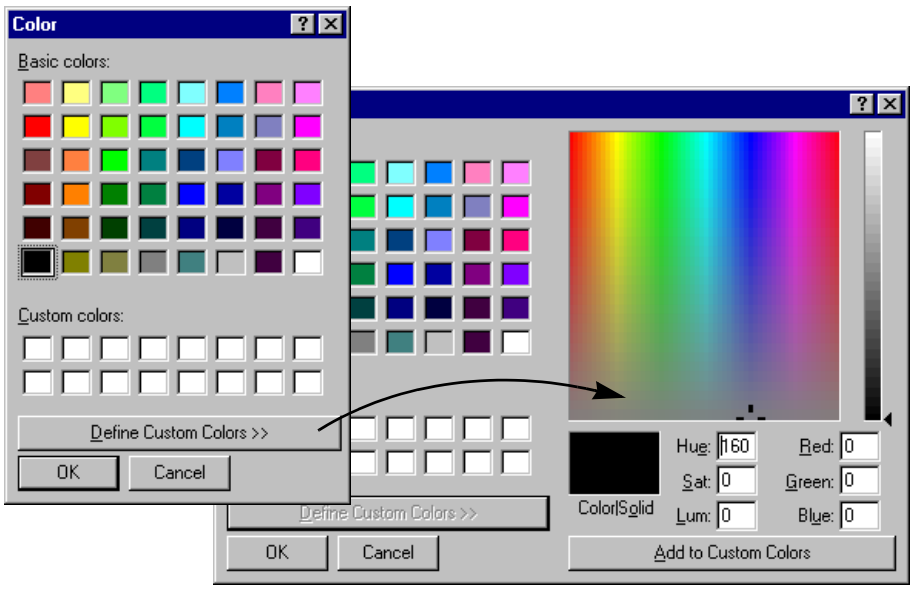

Figure 6-5. Color Window (Expanded with Custom Colors)

c. Select a color to apply to the map:

1.) Click one of the pre-defined **Basic Colors** from the palette.

or

Click **Define Custom Colors** option to expand the window, and use the cross-hair and color slide bar to create a color from the color refiner box. When the color has been defined as you want it, click the **left** mouse button on the **Add to Custom Colors** button to create a new entry in the Custom Colors palate, then click the **left** mouse button to select the custom color you wish to use for the map background.

d. Click **OK** to return to the New Map View window.

2. To add a bitmapped image background to the map:

a. Select the **Bitmap** option.

b. Specify the bitmap file name (and path) in the appropriate text box, or use the **Browse** button to select the file from the appropriate directory.

c. If you want to tile the bitmap (so that it is repeated as many times as necessary to cover the map background), select the **Tile Bitmap** option.

TIP

NetSight Element Manager includes a self-extracting ZIP file that contains a variety of bitmap images you can use for map backgrounds. The file — called "images.exe" — is located in the \NetSight Element Manager\Bin directory; double-clicking on this file will extract the bitmap images.

Setting a Background Grid

You can set up a grid in your map window to assist in drawing map items or in positioning map objects, symbols, and icons. To establish a grid:

1. Click to select the **Display Grid** option to activate the grid display in the map.
2. In the **Grid Size** field, define the size of each grid square by typing in a pixel value.

When you open the Map window, a grid will display in the map background. To turn off the grid, return to the Map View Properties window and click to deselect the **Display Grid** option.

Setting Print Options for the Map

NetSight Element Manager lets you print map window information to any printer configured for your system. If you would like to know how to place map items so they will fall within a page's print area, you can choose to have the pagination appear in your map window. To do so:

1. Click the **Show Page Lines** option. When you open the Map window, dashed yellow and red lines display to indicate the pagination for the map. Each area bounded by dashed lines represents the printable area of a single page.

You can also use the **File —> Print Preview** option from the primary window menu bar to preview how the map information displays on the printed page. Refer to **About NetSight Element Manager** in Chapter 2, **Overview of NetSight Element Manager**, for more information on the Print Preview option.

Setting the Icon Label Font

You can use the **Font** button to select the font type and size that will be used to label your map elements.

1. Click **Font**. The standard Microsoft Font window, Figure 6-6, opens.

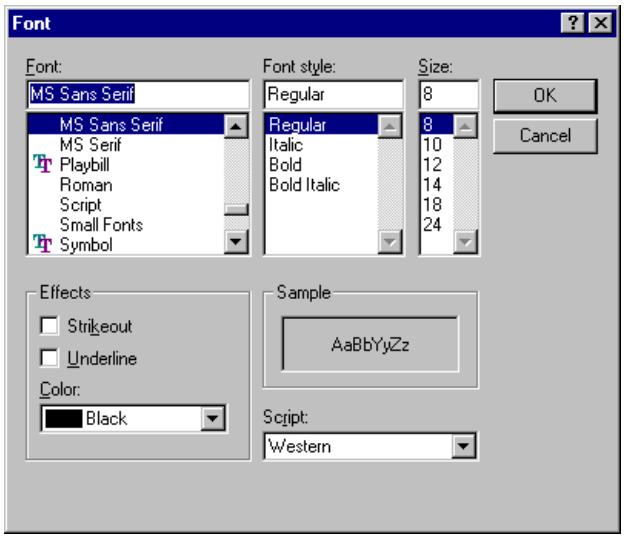

Figure 6-6. The Font Window

2. Select the desired font, style, size, effect, and color using the appropriate fields; note that the effect of the selected options will be displayed in the **Sample** text box.

3. Click **OK** to exit the window and apply your changes. The selected font style will be applied to all icon and symbol labels in the map.

NOTE

Note that the settings in this window only affect the text used in the icon and symbol labels which are applied automatically; to select a font for any added text, use the Text tool described in ***Using Text****, page 6-37.*

Editing Map Properties

You can change any and all of the above settings — including the name of a map or submap — via the Map View Properties window, which is essentially a duplicate of the New Map View window described in the previous section.

To change the properties of an existing map:

1. From the primary window menu bar, with the focus set to the appropriate open map, select **View—>Properties**.

   or

   In the Map View workspace, click the **right** mouse button on the map whose properties you wish to view, then select **Properties** from the resulting menu.

or

From within the map whose properties you wish to view, click the **right** mouse button on any blank space or map element, then select **Map Properties** from the resulting menu.

or

From within the map whose properties you wish to view, double-click anywhere on the map background.

The Map View Properties window, Figure 6-7, opens.

Figure 6-7. The Map Properties Window

Follow the steps outlined in the previous sections to change any of the properties of the selected map.

Adding a Submap

To create a nested hierarchy of maps (as displayed in the Map View workspace), you must add Submap icons to existing maps. You can create a submap by creating an entirely new map, or by linking an existing map; you can also generate submaps automatically (see **Generating Submaps Automatically**, page 6-15). Any changes you make to the nested hierarchy will be reflected in the Map View display. Submap icons both report the status of and provide a means for opening or activating their associated maps.

To add a submap manually:

1. If necessary, open or activate the map to which you wish to add a submap icon.
2. In primary window menu bar, select **Insert—>Sub Map**.

   or

   In the toolbar, click the **Insert Sub Map** icon. The Insert SubMap window, Figure 6-8, displays.

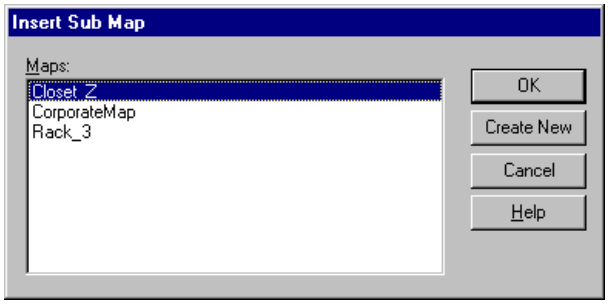

Figure 6-8. The Insert SubMap Window

If no other maps have yet been created, inserting a submap will launch the New Map View window directly.

3. **To create an entirely new map**, click on **Create New**. The New Map View window (as illustrated in Figure 6-4 on page 6-8) opens. Configure the name and display parameters as desired, then click **OK**.

   To link an existing map, select the name of the map you wish to link, then click **OK**. Only the maps available for submap linkage will be listed.

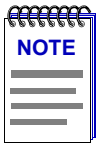

*Remember, any map can only be defined as the submap of **one** root map. If you select a map which is already a submap of another map, the original submap link will be broken, and the new link established. This change will be reflected in the Map View workspace, and the invalid submap icon will be automatically removed from the original map, but you will not be notified in any other way that a submap relationship has been changed.*

4. **If you have chosen to create a new map**, closing the New Map View window changes the arrow cursor to a crosshair cursor; click this cursor in your root map to place the submap icon.

   If you have selected an existing map, it will be opened and placed behind the active map, and the arrow cursor will change to a crosshair cursor. Click this cursor in your root map to place the submap icon.

The Map View workspace will update automatically to display the new map in its position in the hierarchy. The Submap icon itself will display the current status of its associated map; double-clicking on a Submap icon will also open or activate the associated map.

TIP

You can only place a Submap symbol in the root map that was active when the Insert Submap process was initiated.

Generating Submaps Automatically

If you prefer, NetSight Element Manager can create a map hierarchy for you automatically, using the both the subnets you have defined via the Discover Manager and those which have been automatically defined when you added devices to your central node database.

TIP

*For more information on using subnets as part of a discover, see **Defining and Selecting Subnets** in Chapter 3, **Discovering Nodes**. For more information on viewing the subnets that are currently defined, see **Using the Subnets Window**, also in **Chapter 3**.*

Once you have added your network devices to the central node database, you can automatically create a map hierarchy as follows:

1. Select **File—>New—>Subnet Mapping** from the primary window menu bar. Note that this menu option is always available, even when no map windows are open or active. The Subnet Mapping window, Figure 6-9, opens.

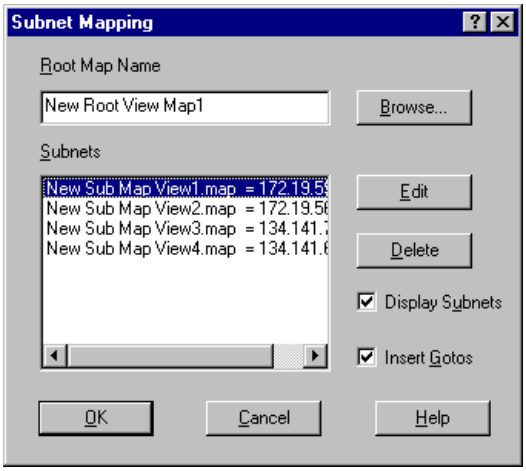

Figure 6-9. The Subnet Mapping Window

2. In the **Root Map Name** field, enter the name you'd like to assign to the root map of your hierarchy. This will be the map that contains the submap icons for each submap that is created; it can be an entirely new map, or you can enter (or use the **Browse** button to locate) an existing map.

3. The **Subnets** list box lists each of the submaps that will be created, along with the default name that has been assigned to each, and the subnet whose devices it will contain. You can edit this default setup as follows:

   a. **To assign map names** more specific than the default name, click to select the submap you wish to change, then click on **Edit**. A window like the one illustrated in Figure 6-10 opens; enter a new name in the text field, then click **OK**. The .MAP extension will be added automatically.

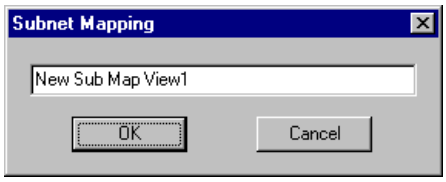

Figure 6-10. Renaming a Submap

   b. **To remove a submap** from the list, click to select the appropriate entry in the list box, then click on **Delete**. When you remove a submap from the list, the submap for that subnet is not generated and the devices in the associated subnet are not displayed in the automatically-generated map hierarchy (though they will still be present in the central node database).

4. Click to select the **Display Subnets** option if you want NetSight Element Manager to place a descriptive text box in each map it creates; this text box will contain the IP designation of the related subnet. This label exists as a simple text element and can be deleted at any time. To leave this text label out, click to de-select the field.

5. Click to select the **Insert GoTos** option if you want NetSight Element Manager to include a GoTo symbol representing the root map in every submap it creates. These symbols can also be added or deleted at a later time, if you choose. If you wish to leave these symbols out, click to de-select this option.

6. When your hierarchy has been configured as desired, click **OK** to begin the submap creation process.

Once they have been created, each of your new maps will be displayed in the primary window, as illustrated in Figure 6-11; they will also be displayed in the Map View workspace. By default, device icons within each submap will be arranged in a grid pattern; if the related options have been selected, a subnet label and GoTo symbol will be displayed at the top of each map window.

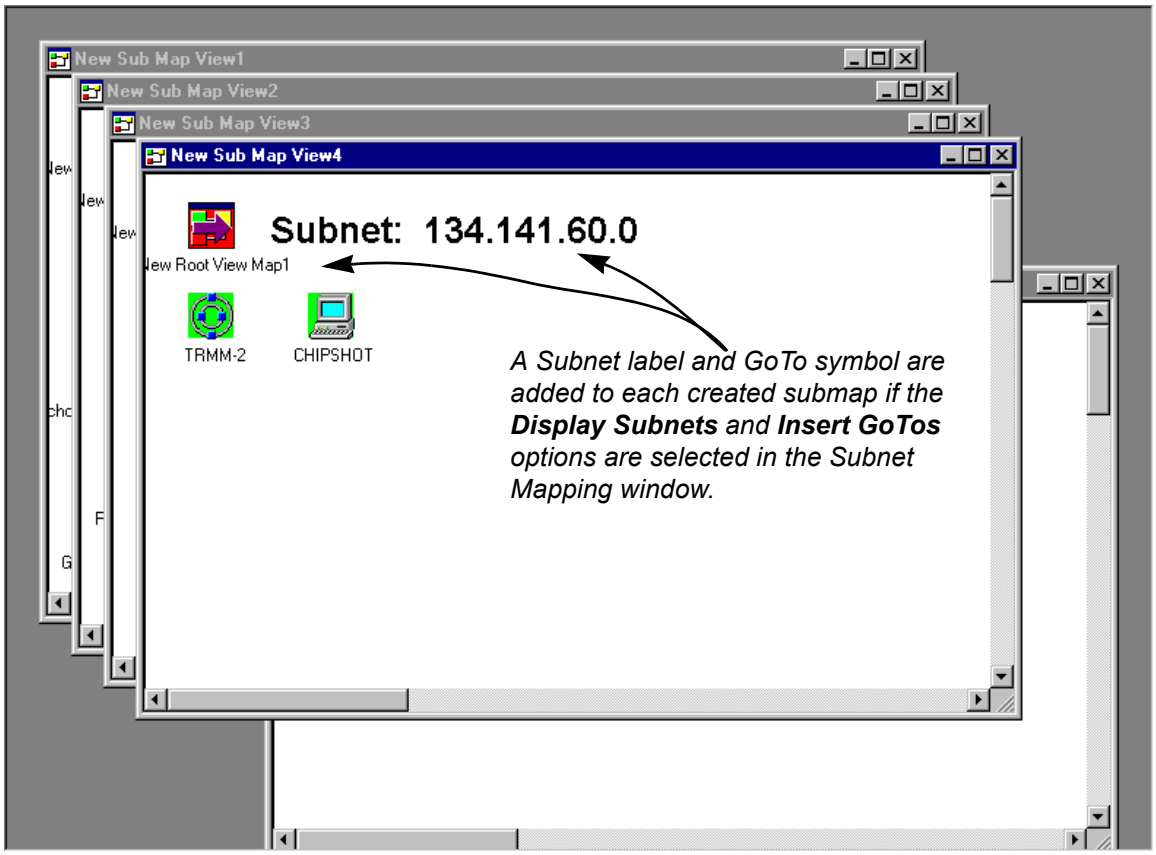
New Sub Map View1
New Sub Map View2
New Sub Map View3
New Sub Map View4
Subnet: 134.141.60.0
New Root View Map1
TRMM-2
CHIPSHOT
A Subnet label and GoTo symbol are added to each created submap if the Display Subnets and Insert GoTos options are selected in the Subnet Mapping window.

Figure 6-11. Automatically Created Submaps

Changing a Submap Link

Whether you have chosen to create a new map, link an existing map, or create a submap hierarchy automatically, submap links are entirely flexible and can be changed at any time. You can change submap relationships by moving a submap icon from one map to another (via drag-and-drop or cut and paste), by placing a submap icon in a new map, or by deleting a submap icon altogether; you can also change the name assigned to a submap without breaking any existing submap links. Just remember the one restriction applied to submap relationships: no map can be a submap to more than one root map (and if you try this, the submap relationship will be changed without warning). All changes will be reflected in the Map View workspace.

*If you use the **Cut** command to remove a submap icon, only the icon will be removed, and no confirmation is required. If you use the **Delete** command, the icon will be removed, and a prompt window opens asking if you wish to remove the view as well. Click **Yes** to remove both the submap icon and the map view it represents; click **No** to remove the icon only.*

*If you accidentally select **Yes**, don't forget you can still recover your map view by importing its associated .MRM file, which will not be removed with the .MAP version. See **Importing an MRM Map**, page 6-20, for more information.*

Adding a GoTo Symbol

You can add GoTo symbols to any map at any time, creating the ability to open or activate one map from another. While the status of the map represented by the GoTo symbol is displayed in the symbol itself, that status is *not* included in the overall status of the map of which the GoTo symbol is a part, and will not be propagated up the map hierarchy. Relationships formed by GoTo symbols are also not displayed in the Map View tree hierarchy.

1. If necessary, open or activate the map to which you wish to add a GoTo symbol.
2. In the primary window menu bar, select **Insert—>GoTo Map**.

   or

   In the toolbar, click the **Insert GoTo Map** icon.

   The Insert Map GoTo window, Figure 6-12, opens.

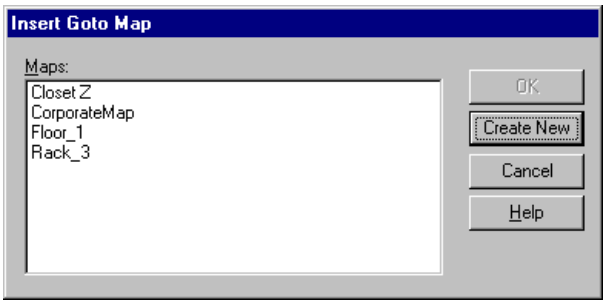

Figure 6-12. The Insert Map GoTo Window

If no other maps have yet been created, inserting a GoTo symbol will launch the New Map View window directly.

3. **To create an entirely new map**, click **Create New**. The New Map View window (as illustrated in Figure 6-4 on page 6-8) opens. Configure the name and display parameters as desired, then click **OK**.

   To link an existing map, select the name of the map you wish to link, then click **OK**. Only the maps available for GoTo linkage will be listed.

4. **If you have chosen to create a new map**, closing the New Map View window changes the arrow cursor to a crosshair cursor; click this cursor in your root map to place the submap icon.

   If you have selected an existing map, it will be opened and placed behind the active map, and the arrow cursor will change to a crosshair cursor. Click this cursor in your root map to place the GoTo icon.

You can only place a GoTo symbol in the root map that was active when the Insert GoTo process was initiated.

Be sure to save the map to which you have added the GoTo symbol; see **Saving Maps**, page 6-62, for more information.

Importing an MRM Map

If you are upgrading from an earlier version of NetSight Element Manager, your map files should be backed up and restored automatically as part of the install process, and should open in their original condition.

Although the .MRM maps are no longer a necessary step in converting your files from one version to another, you may still find that you occasionally need to import an .MRM file. Each map file in NetSight Element Manager is saved in both .MAP format (the format read by NetSight Element Manager) and in .MRM format (a text-based human readable version). If a .MAP file becomes corrupted, you may be able to recover your map by importing the .MRM file. If you inadvertently delete a .MAP file, the accompanying .MRM file is not automatically removed, and it can be re-imported. Or if you want to move a map file from one workspace folder to another, you can achieve this by deleting the .MAP file from the workspace, selecting the folder you want to move the map to, and importing the associated .MRM file into that new folder.

When you import an .MRM file, NetSight Element Manager first checks your current central node database for information about the nodes contained in the .MRM map. If those nodes have already been added to the node database, the appropriate icons are automatically added to the imported map. If those nodes are not yet in the database, NetSight Element Manager will try to re-discover them either individually or as a group, depending on the import option you have selected.

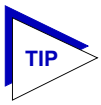

If you will be importing .MRM files which contain many devices not currently present in your central node database, the import process can take a while to complete, as NetSight Element Manager attempts to re-discover each device. To reduce the amount of time it takes to import a map, try adjusting your polling and retry settings via the ***Node*** *page in the* ***Tools—>Options*** *window available from the primary window menu bar.*

If you import an .MRM file which contains a device that is currently in the Excluded Nodes database, the icon will be imported as a gray question mark. Double-clicking on this icon will initiate the normal add device process; however, before the device is added to the map, you will be notified that it is in the Excluded Nodes list, and asked if you want to include it. If you say no, the icon will remain gray; if you say yes, the icon will be displayed normally, and the device will be re-added to the central node database.

If you import a root map, its associated submaps will be imported automatically, and all hierarchical relationships will be restored.

Setting MRM Import Options

The Importing MRM Maps Options window allows you to determine how NetSight Element Manager will handle the re-discovery of nodes in your .MRM maps that are not present in your current central node database.

1. In the primary window menu bar, select **File—>MRM Maps—>Options**; the Importing MRM Maps Options window, Figure 6-13, opens.

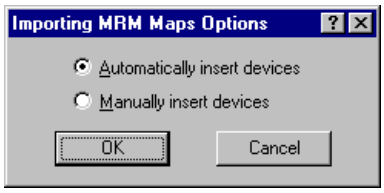

Figure 6-13. Importing MRM Maps Options

2. **If you want all unknown nodes to be discovered at once** (without requiring any input from you), select **Automatically Insert Devices**.

   If you want to be prompted with an Insert Device window for each unknown node, select **Manually Insert Devices**.

3. Click **OK**. The configured option will be used for all .MRM imports until a new option is selected.

Importing a Map

1. In the Map View window, select the folder into which you wish to import the map. If you want to import the map to the top level of the Map View hierarchy, select the main Map Views folder, or highlight any existing map.
2. From the primary window menu bar, select **File—>MRM Maps—>Import MRM Map**. The Import MRM window, Figure 6-14, opens.

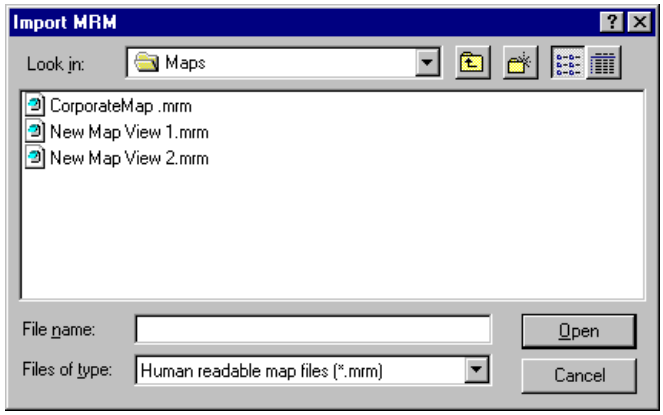

Figure 6-14. The Import MRM Window

3. Use the **Look in** field to select the appropriate directory, then select the file you wish to import in the list box. The selected file opens in the **File name:** text box; click on **Open** to import the file. The Import Progress window, Figure 6-15, opens.

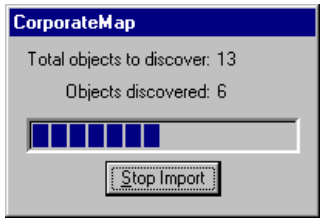

Figure 6-15. MRM Import Progress Indicator

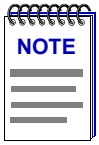

If you import an .MRM map which has the same name as an existing map, one of two things will happen: if the existing map is open, you will be notified that a map of the same name exists and is open, and asked if you want to activate it; if the existing map is ***not*** *open, you will be notified that a map of the same name exists and asked if you wish to overwrite it. You will not have the option of overwriting an open map.*

4. **If you have selected the automatic insertion option**, a mini-discovery will be performed to contact any devices not already present in the central node database. This is the process that can take some time; a message window will be displayed while the discovery is in progress.

   If you have selected the manual insertion option, for each device not located in the central node database you will be prompted with the standard Insert Device window. Use this window to confirm the IP address and community name to be used in the discover process, or click **Cancel** to remove the noted device from the import.

   To stop an import in progress, select **Stop Import**. A message window opens indicating that the import has been stopped. Click **OK** to clear the message; note that the partially-imported map remains in place.

5. When the import is complete, a message window opens noting that fact and indicating the number of nodes which have been added; click **OK** to clear the window. The new map file will be open and active.

Importing a root map will automatically import all associated submaps (if necessary), and re-establish their hierarchical structure.

Adding Network Elements to your Map

For the most part, adding a device icon to a map simply involves dragging the selected icon or icons from any one of the available views or from another map, or using the standard Copy and Paste commands. When you perform a network Discover operation, all discovered devices are automatically added to the central node database maintained by NetSight Element Manager, and they will automatically appear in the appropriate list and tree views; adding individual devices from the primary window menu bar (via **Insert—>Device**) or from the toolbar (via the Insert Device icon) when no map window is active also places the information directly into the central node database (and the appropriate views).

To move a single device from a view onto a map, use the **left** mouse button to click on the device you wish to add, and drag it into the appropriate map window; to move more than one device, use *shift-click* or *ctrl-click* to select all of the nodes you want to move.

However, you can also add a new device directly to a map (and automatically to the node database and all appropriate views); this process is described in the following sections.

Adding a Device Directly to a Map

1. Open or activate (as appropriate) the map window to which you wish to add a new device.
2. In the primary window menu bar, select **Insert—>Device**.

   or

   In the toolbar, click the **Insert Device** icon.

   or

   Anywhere within the selected map window (even on an existing icon, if you want), click the **right** mouse button to display the Map menu (as illustrated in Figure 6-16). Click again to select **New Device**.

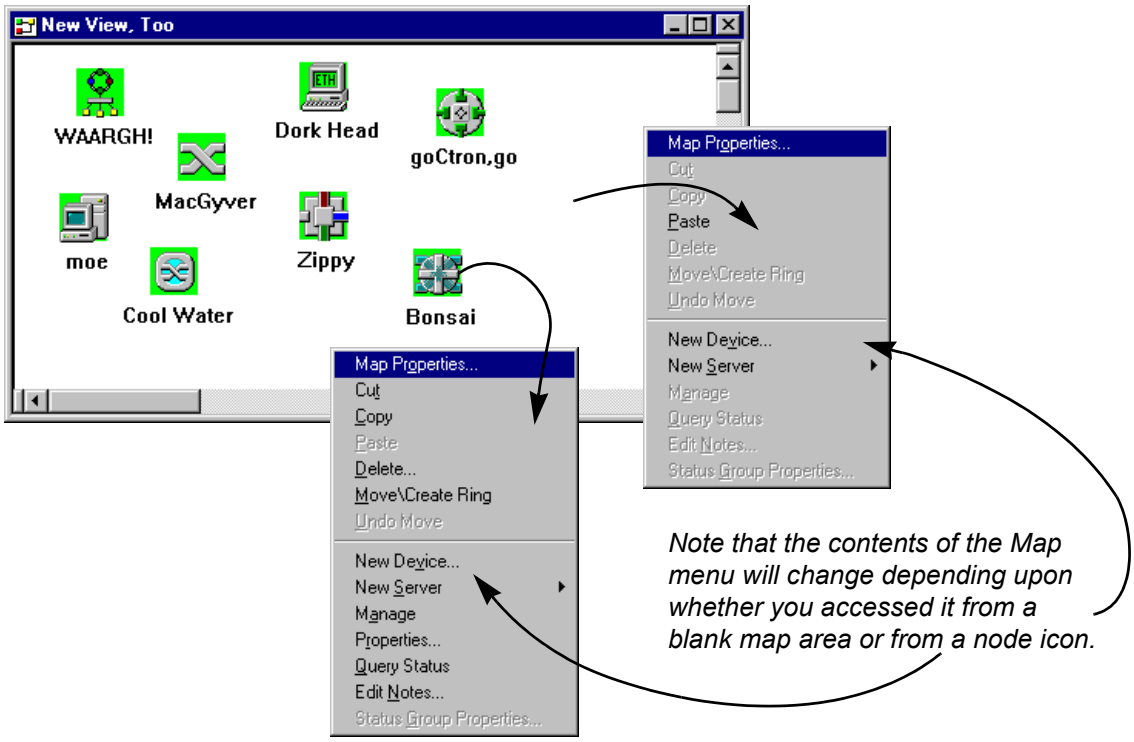

Figure 6-16. Displaying the Map Menu

The Insert Device window, Figure 6-17, opens.

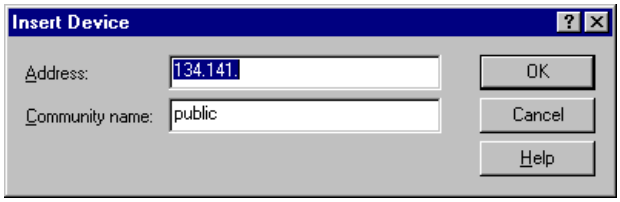

Figure 6-17. The Insert Device Window

3. In the **Address** field, enter the IP address assigned to the device you wish to add.
4. In the **Community Name** field, enter a community name that will provide at least Read access to the device you are adding. For full management access, use a community name that will provide superuser access.

TIP

*If you have configured and activated the Autofill options (either via the Welcome Wizard or via the **Tools—>Options** window), the autofill values you have set for IP address and community name display automatically in the appropriate fields. You can keep these options or replace them, as necessary.*

5. **If the device is successfully contacted**, the Device Properties window, Figure 6-19 (page 6-26), opens. Make any changes to the properties, as desired, then click on **OK**. The arrow pointer will become a crosshair cursor; click the cursor in the active map window to place the new icon.

   If the device cannot be contacted, the Insert Device window will expand to display the **Timeout** and **Retry** counters, and a message window opens (as illustrated in Figure 6-18).

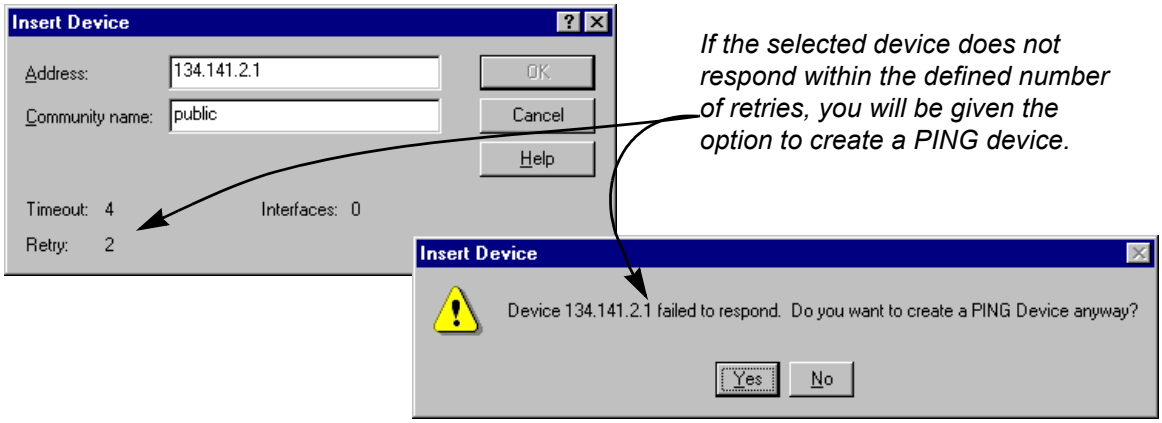

Figure 6-18. Device Timeout Windows

If you click **Yes**, the Device Properties window opens, and you can add an **Unnamed PING device** to your map. If you click **No**, no new device will be added.

TIP

*You can change the icon assigned to an Unnamed PING Device so that it more accurately reflects the actual device being monitored; see **Editing Device Properties** in Chapter 4, **List Views**, for details.*

If the device you are adding is already part of the database, a message window displays after you close the Properties window, indicating that the selected node already exists. Click **OK**, and a crosshair cursor displays allowing you to place the icon; there is no restriction on placing duplicate copies of the same icon in a single map or on multiple maps.

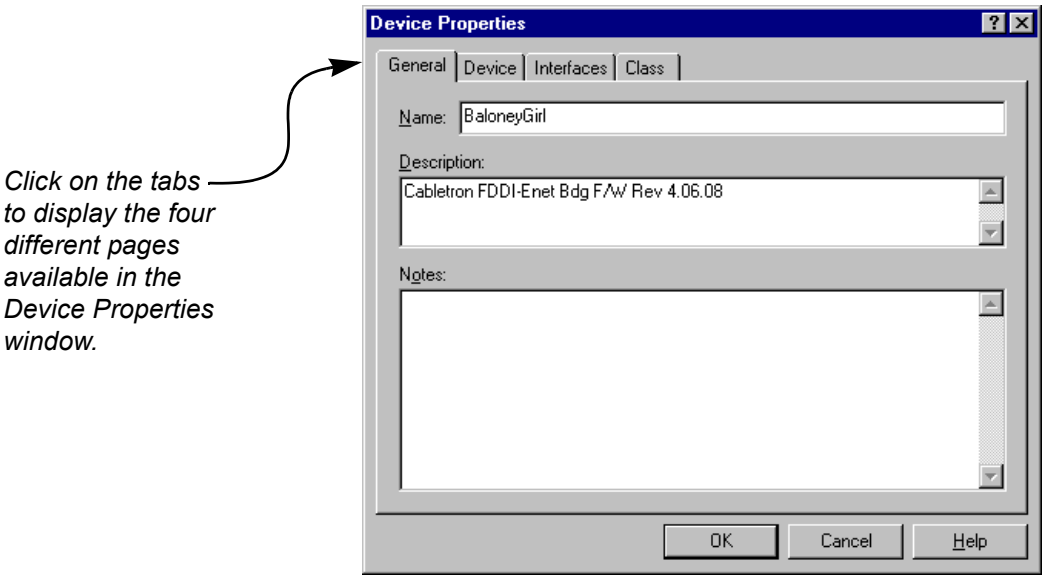

Figure 6-19. The Device Properties Window

By default, device icons will be labeled by the MIB-II *sysName* (if one has been assigned) or, if no name has been assigned, by IP address. You can change the icon name via the Device Properties window; for more detailed, see **Editing Device Properties** in Chapter 4, **List Views**.

Customizing Your Map

In addition to the manageable nodes described in the sections above, your maps can also contain a variety of non-manageable graphical objects that can help you create a better visual display of your network configuration. Alignment and layering tools are also provided that allow you to organize and align map objects with the click of a mouse button. Drawing, layering, and alignment tools are available either from the toolbar (see Figure 6-20) or from the **Draw** and **Layout** menus on the primary window menu bar.

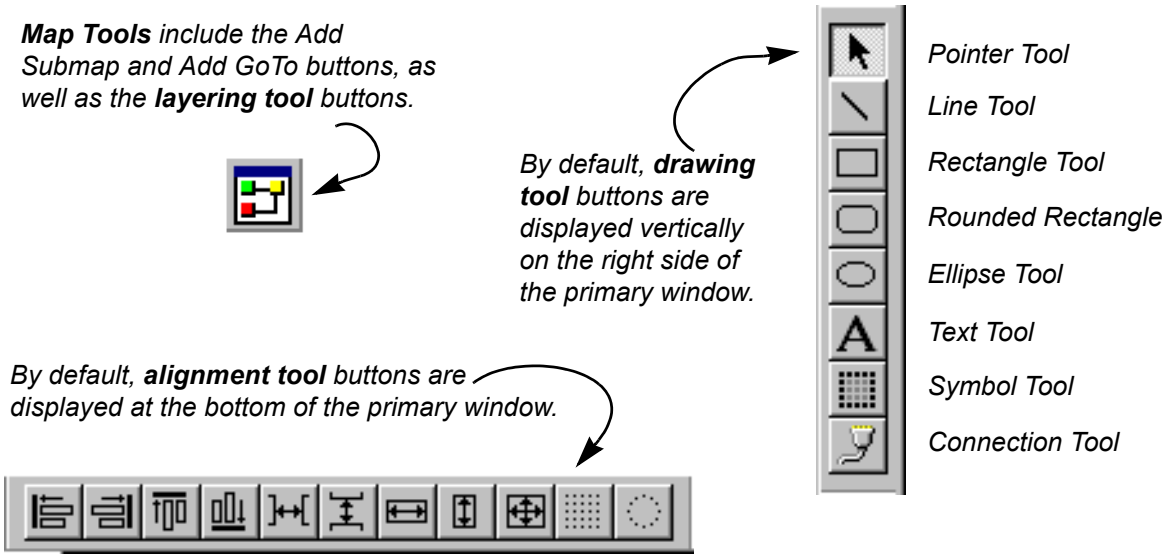
Map Tools include the Add Submap and Add GoTo buttons, as well as the **layering tool** buttons.
By default, **drawing tool** buttons are displayed vertically on the right side of the primary window.
Pointer Tool
Line Tool
Rectangle Tool
Rounded Rectangle
Ellipse Tool
Text Tool
Symbol Tool
Connection Tool
By default, **alignment tool** buttons are displayed at the bottom of the primary window.

Figure 6-20. Drawing, Alignment, and Layering Tools via the Toolbar

You can also annotate any object in your maps by adding an Object Note; the following sections describe how to add non-management graphic objects to your map, as well as how to add an Object Note to any object.

Using Non-manageable Symbols

In addition to the device icons described in the previous sections, you can also add symbol icons to your map to represent network devices you cannot — or choose not to — manage via NetSight Element Manager. A library of predefined symbols are provided, or you can define your own symbols by adding additional bitmap images.

Symbols — even those which look just like device icons — do not report status and do not provide access to management.

Placing a Symbol Icon

To add an existing symbol to your map:

1. If necessary, open or activate the map to which you wish to add a symbol. None of the Draw toolbar options will be available unless at least one map is open and active.
2. Click the Symbol icon from the toolbar.

   or

   Select **Insert—>Symbol** from the primary window menu bar.

The Select Symbol window, Figure 6-21, opens.

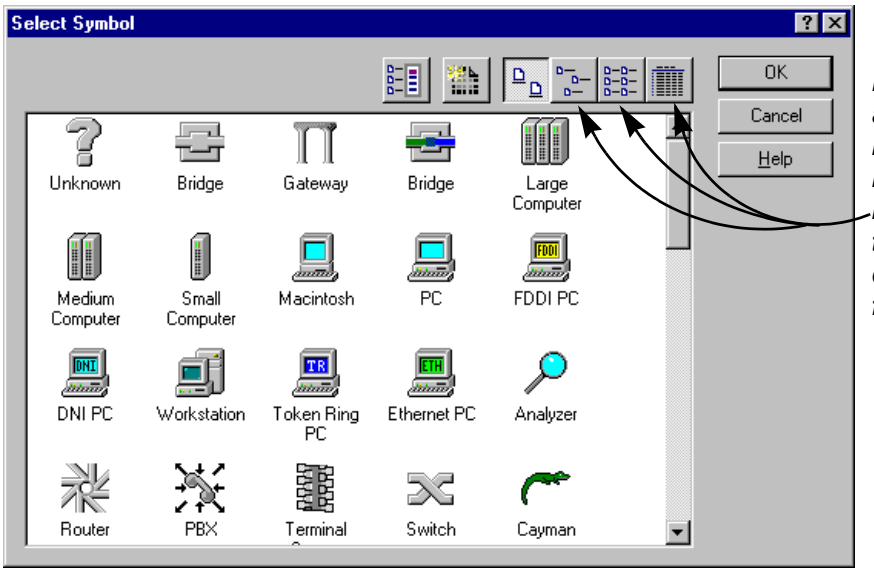

*By default, the available symbols will be displayed in **Large Icon** format; use the buttons at the top of the window to select a different display format, if desired.*

Figure 6-21. Select Symbol Window

3. Select the symbol you wish to add. By default, the available symbols are displayed in Large Icon format; use the buttons at the top of the window (see Figure 6-21) to change the display, if desired, or use the scroll bar to view all available symbols. Once you have selected your symbol icon, click **OK**.
4. Move the cursor back into the map window, and note that it has changed from an arrow to a crosshair. Position the crosshair where you want to place the symbol, then click to place the symbol in your map. The symbol will be displayed with its default label. (To change the label, see **Editing Individual Symbol Properties**, page 6-29.)

These symbols do not report node status nor provide access to node management.

Editing Symbol Properties

If you wish, you can edit symbol properties on two levels: you can edit the properties of an *individual* symbol you have placed on a map, assigning it a unique label; you can also edit the symbol *database*, either by changing the way a symbol is defined or displayed, or by creating an entirely new symbol.

When you edit the symbol database, you are changing the basic information NetSight Element Manager uses to define and display symbols. These changes cannot easily be undone; if you choose to do this, exercise appropriate caution.

Editing Individual Symbol Properties

When you place a symbol in your map, it will be labeled with the global label assigned to the symbol type in the symbol database. You can change this label for each individual symbol icon via the Symbol Properties window, Figure 6-22, as follows:

1. To access the Symbol Properties window from the appropriate map window, double-click on the symbol icon.

   or

   Select the symbol, then select **Edit —> Properties** from the primary window menu bar.

   or

   Select the symbol, then select the **Properties** icon from the toolbar.

   or

   Click the **right** mouse button on the symbol, then select **Properties** from the resulting menu.

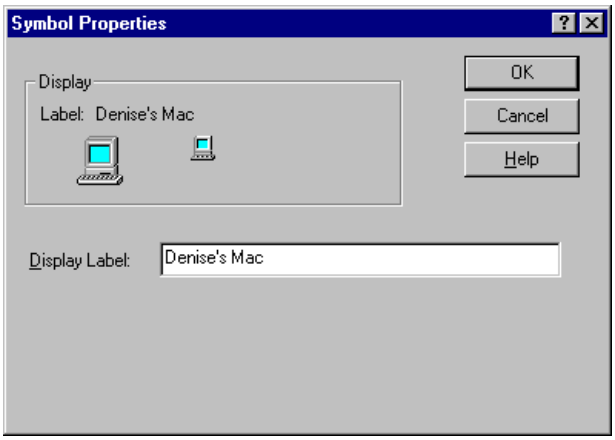

Figure 6-22. Editing Symbol Properties

2. In the **Display Label** field, enter the new label you wish to assign to the *selected symbol icon only*, then click **OK**.

The new label will be displayed on the selected symbol icon.

Editing the Symbol Database

If you wish, you can also edit the properties assigned to a symbol type in the NetSight Element Manager symbol database. If you choose to edit a symbol at this level, however, note that the changes you make cannot be easily undone, and will be applied to every instance of the symbol that you place in your maps.

TIP

If you can't find a symbol in the database that suits your needs, but you are reluctant to permanently edit the symbol database, you might want to create an entirely new symbol instead. This new symbol will be added to the existing database, but will not change any existing symbol definitions. See ***Creating a New Symbol****, page 6-32, for details.*

To edit the symbol database:

1. If necessary, open or activate a map window. (Remember, none of the Draw toolbar options will be available unless at least one map is open and active.)
2. Click the Symbol icon from the toolbar.

   or

   Select **Insert—>Symbol** from the primary window menu bar.

   The Select Symbol window, Figure 6-23, opens.

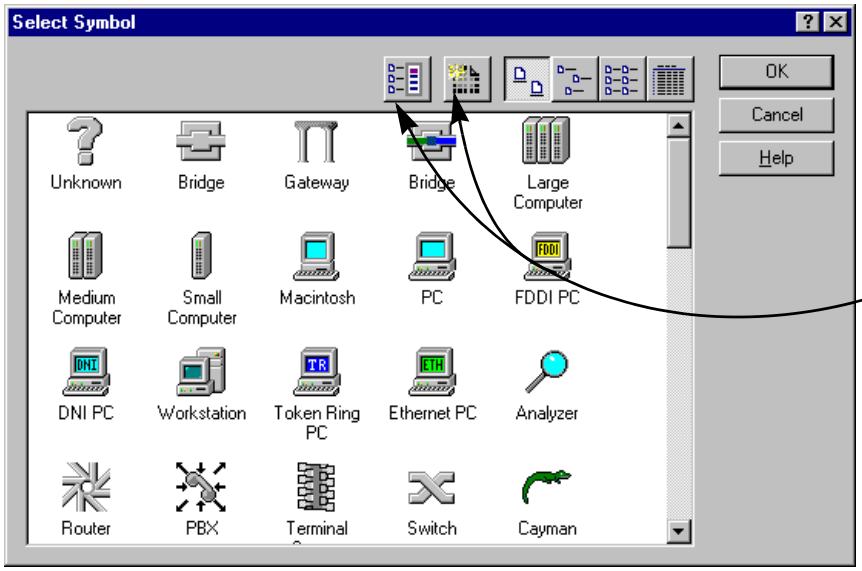

Figure 6-23. Select Symbol Window

3. Select the symbol you wish to edit. By default, the available symbols are displayed in Large Icon format; use the buttons at the top of the window to change the display, if desired, or use the scroll bar to view all available symbols. Once you have selected your symbol icon, click the Properties button. The symbol Properties window, Figure 6-24, opens.

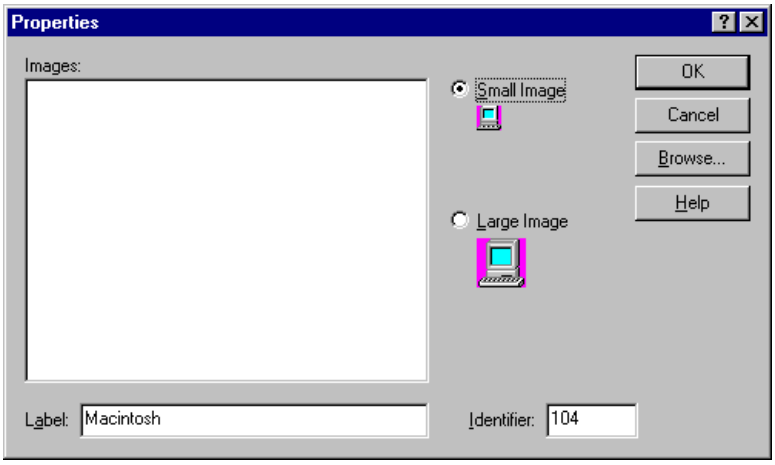

Figure 6-24. Global Symbol Properties Window

4. To change the label assigned to all symbols of the selected type, click in the **Label** field, and enter the desired text.

5. To change the numeric identifier assigned to the symbol type, click in the **Identifier** field, and enter a new numeric value.

If you do choose to edit the Identifier value, be sure you do not assign a value already being used by another symbol; if you do, the symbol you are editing will ***replace*** *the symbol originally identified by that value. There is no functional reason to edit this value; we recommend that you do not.*

6. To edit the image displayed in the selected symbol's icon:

   a. Click **Browse**; in the resulting File Open window, select the bitmap file you wish to use as a basis for the symbol image, then click **Open**. The selected bitmap image will be displayed in the **Images:** display area in the Properties window. (If the image is too big in either dimension to be fully displayed, a message window displays, indicating that the image will be cropped to fit.)

   b. Click the **Small Image** radio button, then move the resulting square selection cursor around on the image to select the portion you wish to display for a small (16 x 16 pixels) symbol icon, and click again to make the selection. The image portion you have selected will be displayed below the **Small Image** radio button. Move the cursor and click again to change this selection until you find the one you want.

   c. Click the **Large Image** radio button, then move the resulting square selection cursor around to select the image portion you wish to display for a large (32 x 32 pixels) symbol icon, and click again to make the

selection. The image portion you have selected will be displayed below the **Large Image** radio button. Move the cursor and click again to change this selection until you find the one you want.

7. Click **OK** to save your new symbol definition.

If you have changed the symbol label or image, the new properties will be displayed in the Select Symbol window.

Creating a New Symbol

If you can't find an existing symbol that meets your needs, and you are reluctant to edit the existing symbol database, you can create your own new symbol, as follows:

1. If necessary, open or activate a map window. (Remember, none of the Draw toolbar options will be available unless at least one map is open and active.)
2. Click the Symbol icon from the toolbar.

   or

   Select **Insert—>Symbol** from the primary window menu bar.

   The Select Symbol window, Figure 6-23 (page 6-30), opens.
3. Click the New button, or **right-click** on any existing symbol, and select **New** from the resulting menu. The New Symbol window, Figure 6-25, opens. (Note that this window is virtually identical to the symbol Properties window described in the previous section.)

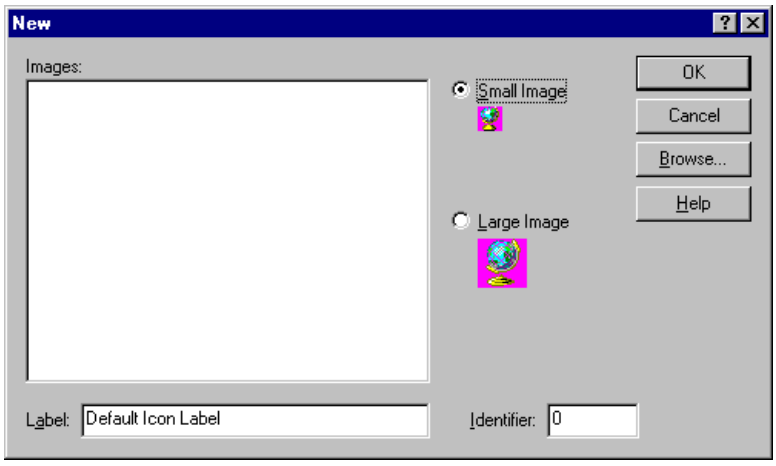

Figure 6-25. New Symbol Window

4. To assign a label for your new symbol, click in the **Label** field, and enter the desired text.

5. To assign a numeric identifier to your new symbol, click in the **Identifier** field, and enter a new numeric value.

*Be sure you do not assign a value already being used by another symbol; if you do, the symbol you are creating will **replace** the symbol originally identified by that value! To verify that you are not using an existing identifier, view the Select Symbol window (Figure 6-23, page 6-30) in Small Icon Detail format by selecting the* *button; the identifiers assigned to each existing symbol will be displayed, and can be sorted for easier viewing.*

6. To edit the image that will be displayed in your new symbol's icon:

   a. Click the **Browse** button; in the resulting File Open window, select the bitmap file you wish to use as a basis for the symbol image, then click **Open**. The selected bitmap image will be displayed in the **Images:** display area in the Properties window. (If the image is too big in either dimension to be fully displayed, a message window opens, indicating that the image will be cropped to fit.)

   b. Click the **Small Image** radio button, then move the resulting square selection cursor around on the image to select the portion you wish to display for a small (16 x 16 pixels) symbol icon, and click again to make the selection. The image portion you have selected will be displayed below the **Small Image** radio button. Move the cursor and click again to change this selection until you find the one you want.

   c. Click the **Large Image** radio button, then move the resulting square selection cursor around to select the image portion you wish to display for a large (32 x 32 pixels) symbol icon, and click again to make the selection. The image portion you have selected will be displayed below the **Large Image** radio button. Move the cursor and click again to change this selection until you find the one you want.

7. Click **OK** to save your new symbol definition.

Your new symbol will be displayed in the Select Symbol window.

Using Connections

To symbolically represent a network connection, you can create a link between any two symbols or node icons. This link appears as a dynamic line between the two selected symbols and/or node icons that will redraw itself automatically when one or the other connected object is moved.

To draw a connection between two nodes or symbols:

1. If necessary, open or activate the map to which you wish to add a connection. Note that none of the toolbar or menu drawing options will be available unless at least one map is open and active.

2. Click the Connection tool from the toolbar (see Figure 6-20, page 6-27).

   or

   Select **Draw—>Connection** from the primary window menu bar.

3. Move the cursor back into the map window, and note that it has changed from an arrow to a ⊘ symbol. This symbol will change to a crosshair when the mouse is positioned over an object (node icon or symbol) which can support a connection. Position the crosshair over the first of the two elements you wish to connect, click the mouse, and release.

4. Click in any blank map area to add segments, or vertices, to your connection. One additional line segment will be added each time you click the mouse in a blank map area.

5. To complete the connection, move the cursor to the second of the two elements you wish to connect, and click the mouse button again. Once the connection is complete, the cursor will regain its arrow state; to add additional connections, you must select the Connection tool again.

Once you have added your connection, you can move either one of the two connected symbols, and the connection will be re-drawn as appropriate. You can also add multiple connection lines to a single symbol, and all will be maintained and re-drawn as necessary.

You can also edit the size, type, and color properties of the connection line, just as you would edit a regular line; see page 6-35 for more information on accessing the Line Properties window and editing line properties.

Using Lines and Shapes

The Line, Rectangle, Round Rectangle, and Ellipse tools allow you to add customizable shapes to your map.

1. If necessary, open or activate the map to which you wish to add a graphic element. Note that none of the toolbar or menu drawing options will be available unless at least one map is open and active.

2. Select the appropriate tool from the toolbar (see Figure 6-20, page 6-27).

   or

   Select the appropriate option from the **Draw** menu in the primary window menu bar.

3. Move the cursor back into the map window, and note that it has changed from an arrow to a crosshair. Position the crosshair where you want to begin drawing.

If you are creating a rectangle or ellipse, click and hold the mouse button and drag to create the size and shape object you want. Releasing the mouse button completes the object, but note that the crosshair cursor remains, allowing you to draw additional objects of the same shape without having to re-select the tool.

If you are creating a line, click the mouse button once to place the beginning line point, then move the cursor in the direction you wish to draw your line. Each time you single-click on the mouse button, a line segment is completed; at the end of each line segment, you can either change directions and draw an additional, attached line segment, or double-click to complete the line. Note that once your line is complete, the crosshair cursor is returned to its pointer state; to draw additional lines, re-select the line tool.

Once you've created an object, you can edit its graphical properties (such as fill, line type, and color) via the appropriate Properties window.

To access the Properties window:

1. Double-click on the object.

   or

   Select the object and then select **Edit —> Properties** from the primary window menu bar.

   or

   Select the object, then select the **Properties** icon from the toolbar.

   or

   Right-click on the object, then select **Properties** from the resulting menu.

   The appropriate Properties window, as illustrated in Figure 6-26, opens.

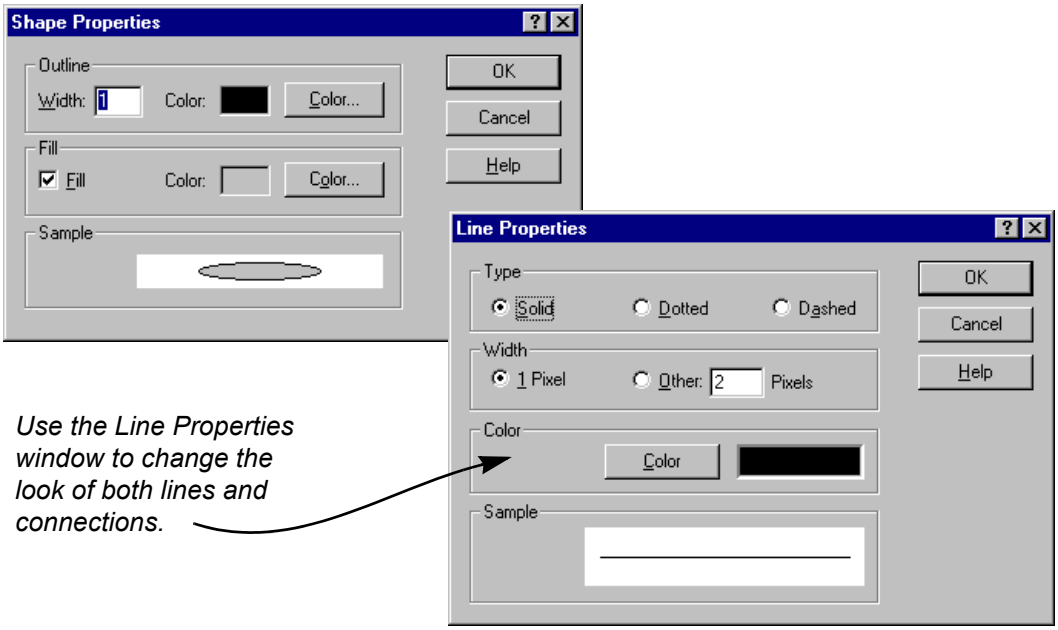

Figure 6-26. Line and Shape Properties

To specify Line Properties:

1. Select the **Line Type** by clicking the mouse button on the appropriate option: **Solid**, **Dotted**, or **Dashed**.
2. If you selected a Solid line, you can specify its **Line Width** in pixels. Dotted and dashed lines are fixed at 1 pixel.
3. Select the **Line Color** by clicking on **Color** to bring up the Color Palette window, selecting a pre-defined **Basic Color** or a defined **Custom Color** (see Figure 6-5, page 6-10), and clicking **OK** to return to the Line Properties window.
4. The Sample portion of the window will display your changes as you make them; when the sample displays the look you want to achieve, click **OK** to apply changes and exit the window.

To specify Shape Properties:

1. Specify the **Width** of the shape's outline by typing in a pixel value; specify the **Color** of the outline by using the **Color** button to launch the Color Palette window (see Figure 6-5, page 6-10).
2. To fill the shape, select the **Fill** option; specify the **Color** of the fill by using the **Color** button to launch the Color Palette window.

3. The Sample portion of the window will display your changes as you make them; when the sample displays the look you want to achieve, click **OK** to apply changes and exit the window.

You can also reshape existing objects, as follows.

To reshape an existing line:

1. Click once on the line to display the handles for its vertices.
2. Move the mouse cursor over a handle to activate reshaping for that vertex (the cursor will display as a dual-headed arrow).
3. Click on the handle and drag the vertex to a new position (its neighboring vertices will remain fixed).
4. Release the mouse and move the cursor off the handle to restore the pointer.

To resize an existing shape:

1. Click once on the object to display its resizing handles.
2. Move the mouse cursor over a handle to activate resizing (the cursor will display as a dual-headed arrow).
3. Click on the handle and drag horizontally, vertically, or diagonally to resize the shape.
4. Release the mouse and move the cursor off the handle to restore the pointer.

Using Text

In addition to the labels provided with icon and symbol elements, you can also add independent text boxes to your map.

1. If necessary, open or activate the map to which you wish to add text. Note that none of the toolbar or menu drawing options will be available unless at least one map is open and active.
2. Select the Text tool from the toolbar (see Figure 6-20, page 6-27).

   or

   Select **Draw—>Text** from the primary window menu bar.
3. Move the cursor back into the map window, and note that it has changed from an arrow to an I-bar. Position the cursor where you want to begin typing, and click the mouse button. A text box opens in the map.
4. Enter text into the text box as desired. The text box will automatically extend to the right as you type, but note that you cannot add a carriage return.

5. To move the text box, click and hold the mouse button, and drag the cursor to the appropriate location. (Note that you can move this text element any time, just like any other map element.)
6. When you have finished typing, click anywhere outside the text box.

The default selections for font type, color, and size will be used for each new text element; to change the style or font of the text in the selected text element:

1. Double-click on the text.

   or

   Select the text, then select **Edit —>Properties** from the primary window menu bar.

   or

   Select the text, then select the **Properties** icon from the toolbar.

   or

   Right-click on the text, then select **Properties** from the resulting menu.

   The Font Properties window, Figure 6-27, opens.

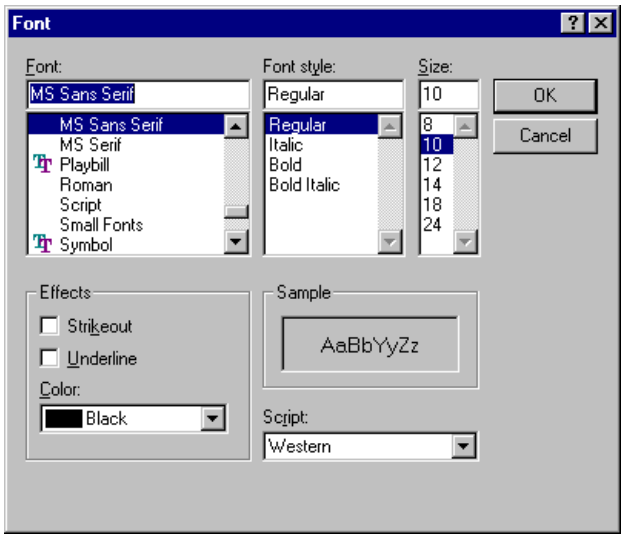

Figure 6-27. Font Properties

2. Using the **Font:** list box, click to highlight one of your installed system fonts to apply to the string.
3. Using the **Font Style:** list box, select a style to apply to the text string: Regular, Italic, Bold, or Bold Italic.

4. Using the **Size:** list box, select a font size for the string.
5. If you want your text to have a **Strikeout** or **Underline Effect**, click to select (X) the appropriate option.
6. Using the **Color** scrollbox, click to select a color for your text string.
7. As you edit the font properties, you can see the results in the **Sample** display. When you have finished, click **OK** to apply changes and exit the Font Properties window.

Note that changes made in this window apply to the entire text string; you cannot apply text effects to individual words or characters.

Adding Object Notes

You can annotate any object in your Map — whether it be a node icon, geometric shape, symbol, text block, connection, or even an embedded or linked object (described in **Using Object Linking and Embedding (OLE)**, beginning on page 6-47) — with object notes; you can also add object notes to nodes which appear in a List or Tree View. Object notes are simply additional descriptions that are associated with the selected element; they have no other function.

You can add an Object Note to any element in a map or other view via the **Edit—>Notes** option on the primary window menu bar, via the **Edit Notes** icon on the toolbar, or via the **Edit Notes** option available on the right-mouse Map or View menu; for device icons, you can also enter a note via the General tabbed page in the Properties window (see Chapter 4, **List Views**, for more information on configuring device properties).

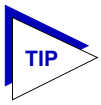

*You can also add notes to multiple nodes via the Multi-Node Properties window; see **Adding Notes to Multiple Nodes**, page 6-42.*

To add or edit an object note to any element:

1. If necessary, open or activate the map or view which contains the object you wish to annotate, then click to select the object.
2. In the primary window menu bar, select **Edit—>Notes**.

   or

   In the toolbar, click **Edit Notes** icon.

   or

   Right-click on the selected element; on the resulting menu, drag down to **Edit Notes**, and release.

   The Object Notes window, Figure 6-28, opens.

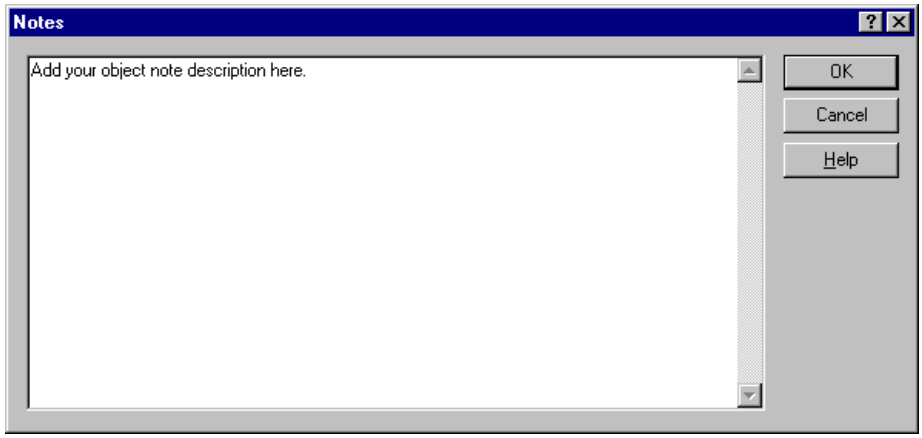

Figure 6-28. Object Notes Window

3. Click to place the cursor in the text area, and add your note. Note that the text will wrap automatically, or you can enter a carriage return by pressing **Enter** (if you are annotating a device icon) or **ctrl-Enter** (if you are annotating a map object).

TIP

*Pressing **Enter** in a Notes window for a map object saves your changes and closes the window.*

4. If you have copied text to the clipboard from another location and wish to place it in the object note, **right-click** to access the text editing menu, and drag down to **Paste**. You can also use this menu to **Copy** text to the clipboard for pasting to another location.
5. When you have finished entering text, click **OK** to save your changes and close the Object Note window.

To add or edit an object note to a device via the Properties window:

1. Open and/or activate the list, tree, or map view which contains the device you wish to annotate, and click on the device icon to select it.
2. In the primary window menu bar, select **Edit—>Properties**.

   or

   In the toolbar, click the **Properties** icon.

   or

   Right-click on the selected device, then select **Properties** from the resulting menu.

3. Enter the desired notes text in the **Notes** field available on the General Properties tabbed page. As illustrated in Figure 6-29, text entered in this Notes field will also be displayed in the Object Notes window, and vice versa.

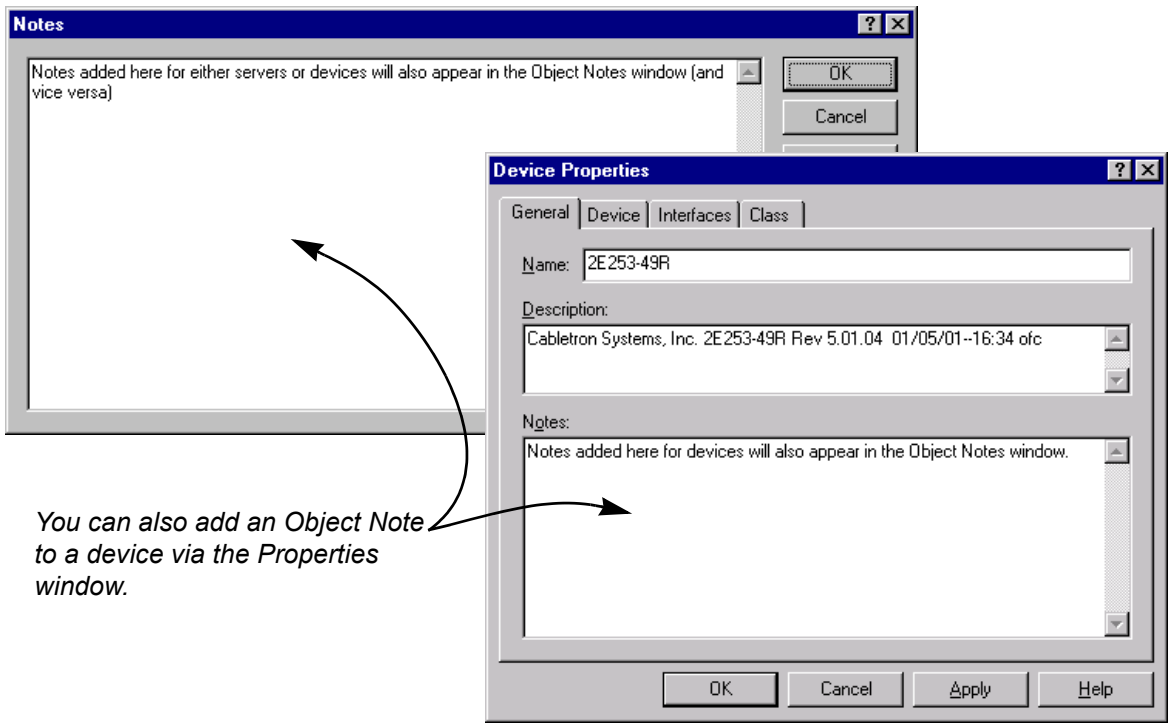

Figure 6-29. Entering an Object Note via the Properties Window

4. Click **Apply** to save your changes (if applicable), then click **OK** to close the Properties window.

TIP

*When an annotated object is duplicated (either in the same map, or in other maps), the notes will be duplicated along with the object. When copying and pasting an object of any kind, you may want to delete or edit the notes associated with that object, to avoid confusion. (Using the **Paste Special (No IPs)** option to duplicate device icons will remove any attached notes along with the other descriptive information.)*

Adding Notes to Multiple Nodes

You can also add the same note to multiple node icons simultaneously via the Multi-Node Properties window. Doing so will overwrite any existing notes attached to the selected nodes, without warning.

To add a single note to multiple selected nodes:

1. Open and activate the view which contains the nodes to which you'd like to add a shared note, and use *shift-click* or *ctrl-click* to select the appropriate nodes. (Or, if you're using a Map View, you can click and drag to select multiple icons.)
2. In the primary window menu bar, select **Edit—>Properties**.

   or

   In the toolbar, click **Properties** icon.

   or

   Right-click on the selected device, then select **Properties** from the resulting menu.
3. Enter the desired notes text in the **Notes** field available on the General Properties tabbed page. As illustrated in Figure 6-29, text entered in this Notes field will also be displayed in the Object Notes window and the Properties window for each individual selected node.

Notes added via the Multi-Node Properties window will overwrite any existing notes on the selected nodes, without warning. Once you have added a note to multiple nodes, it will be displayed in each individual node's Properties and Object Note windows; you cannot display it in the Multi-Node Properties window, however, unless the same group of nodes (or a subset of them) is selected.

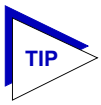

*You can also change the poll settings for multiple nodes; see **Editing Poll Settings for Multiple Nodes** in Chapter 4, **List Views**, for more information.*

Changing the Layout of Map Objects

NetSight Element Manager provides a variety of layout tools that allow you to arrange node icons, symbols, and all other map elements with the click of a mouse. You can layer objects; align a group of selected objects by left edge, right edge, top, or bottom; re-size objects to match the controlling object's height, width, or both; space selected objects evenly (either horizontally or vertically); arrange selected objects in a grid pattern; or even create a ring of objects. Each of these features is described below.

The Layout Tools are available both via the Map and Alignment Toolbars (see Figure 6-20, page 6-27) and the **Layout** menu on the primary window menu bar; these options will only be active when a map window is active.

Layering Items in a Map

You can arrange any map object — node icons, symbols, shapes, lines, connections, and text elements — into layers by using the Layering Tools (see Figure 6-20, page 6-27) or via the **Layout—>Layer** menu on the primary window menu bar. The layering position of a single object is determined relative only to other contiguous objects — not to all objects in the map.

1. If necessary, open or activate the map which contains the elements you wish to layer, then click on the element you wish to move. (Note that you can change the layer for only one element at a time.)
2. Select one of the following toolbar or menu options:

   To move the object to the forefront, or top layer, select the toolbar button or **Layout—>Layer—>Send to Front** on the primary window menu bar.

   To move the object to the background, select the toolbar button or **Layout—>Layer—>Send to Back** on the primary window menu bar.

   To move the object forward one layer, select the toolbar button or **Layout—>Layer—>Step Forward** on the primary window menu bar.

   To move the object back one layer, select the toolbar button or **Layout—>Layer—>Step Backward** on the primary window menu bar.

Align Left, Right, Top, or Bottom

1. If necessary, open or activate the map which contains the objects you wish to align. (Remember, neither the Alignment Toolbar nor the Layout menu options will be active unless a map window is active.)
2. Press and hold the **Ctrl** key, select the objects you wish to align. The last object you select will be the **controlling object** — the object to which the other selected objects will be aligned.

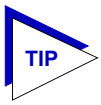

*You can also select objects using the click-and-drag selection method, or via the **Edit—>Select All** options on the primary window menu bar; however, these methods do not allow you to select the controlling object. For the most precise alignment control, the **ctrl-click** method described above is recommended.*

3. Once you've selected the objects you wish to align, release the **Ctrl** key and click on the appropriate alignment tool, as follows:

   To align selected objects to the left edge of the controlling object, click on the Align Left button or select **Layout—>Align—>Left** from the primary window menu bar.

   To align selected objects to the right edge of the controlling object, click on the Align Right button or select **Layout—>Align—>Right**.

To align selected objects to the top of the controlling object, click on the Align Top button or select **Layout—>Align—>Top**.

To align selected objects to the bottom the controlling object, click on the Align Bottom button or select **Layout—>Align—>Bottom**.

4. The aligned objects remain selected after the alignment procedure is complete; you can move the group as a whole by clicking and holding the **left** mouse button, and dragging the group of objects to the desired map location. To de-select the objects, click anywhere in the map background.

Note that these tools align only the selected edge; to control the spaces *between* objects, use the Space Evenly options described below.

Space Evenly (Horizontally or Vertically)

You can also use the layout tools to evenly space selected objects between the left-most and right-most selected objects (**Space Evenly—>Horizontally**) or between the top and bottom selected objects (**Space Evenly—>Vertically**). When you use this feature, the outside objects (left and right, or top and bottom) don't move, so make sure there's enough room between the two outside objects to arrange the remaining selected objects. Note, too, that these tools only control the space *between* objects, and will not arrange the objects in a straight line — for that, you will need to use the Align Top, Bottom, Left, or Right tools described above.

1. If necessary, open or activate the map which contains the objects you wish to space. (Remember, neither the Alignment Toolbar nor the Layout menu options will be active unless a map window is active.)

2. Select the objects you wish to space by clicking the mouse button and dragging to enclose the target objects, pressing and holding the **Ctrl** key and select the target objects, or by using the **Edit—>Select All** or **Edit—>Select All Nodes** options on the primary window menu bar. Remember, spacing will take place between the top and bottom selected objects (for vertical spacing) or between the left-most and right-most objects (for horizontal spacing); these outside objects will not be moved.

3. Once you've selected the objects you wish to space, release the **Ctrl** key (if necessary) and click on the appropriate alignment tool, as follows:

   To space selected objects horizontally, click on the Space Horizontally button or select **Layout—>Space Evenly—>Horizontally** from the primary window menu bar.

   To space selected objects vertically, click on the Space Vertically button or select **Layout—>Space Evenly—>Vertically**.

4. The spaced objects remain selected after the alignment procedure is complete; you can move the group as a whole by clicking and holding the **left** mouse button, and dragging the group of objects to the desired map location,

or align the objects by the top, bottom, left, or right edges using the alignment tools described in the previous section. To de-select the objects, click the **left** mouse button anywhere in the map background.

Make Same Width, Height, or Size

You can use the **Make Same Size** tools to reshape graphic elements in relation to one another. You can resize lines, rectangles, and ovals; text blocks are also technically "resized," but neither the font nor the text's position in the text block are affected. (To change the font, use the Font Properties window, described on page 6-37.) You cannot resize node icons or symbols, nor can icons or symbols serve as the controlling object for graphic element resizing.

1. If necessary, open or activate the map which contains the objects you wish to align. (Remember, neither the Alignment Toolbar nor the Layout menu options will be active unless a map window is active.)
2. Press and hold the **Ctrl** key, and select the object(s) you wish to resize. The last object you select will be the **controlling object** — the object whose width, height, or overall size will be applied to the other selected objects.

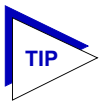

*You can also select objects using the click-and-drag selection method, or via the **Edit—>Select All** or **Edit—>Select All Nodes** options on the primary window menu bar; however, these methods do not allow you to select the controlling object. For the most precise sizing control, the **ctrl-click** method described above is recommended.*

3. Once you've selected the objects you wish to align, release the **Ctrl** key and click on the appropriate sizing tool, as follows:

   To make the selected objects the same width as the controlling object, click on the Make Same Width button or select **Layout—>Make Same Size—>Width** from the primary window menu bar.

   To make the selected objects the same height as the controlling object, click on the Make Same Height button or select **Layout—>Make Same Size—>Height**.

   To make the selected objects the same size (both width and height) as the controlling object, click on the Make Same Size button or select **Layout—>Make Same Size—>Both**.

4. The resized objects remain selected after the alignment procedure is complete; you can move the group as a whole by clicking and holding the mouse button, and dragging the group of objects to the desired map location. To de-select the objects, click in the map background.

Align Objects in Grid

The Grid tool allows you to align all selected objects into a grid formation starting in the upper-left-hand corner of the map window.

1. If necessary, open or activate the map which contains the objects you wish to align. (Remember, neither the Alignment Toolbar nor the Layout menu options will be active unless a map window is active.)
2. Press and hold the **Ctrl** key, and select the object(s) you wish to align in a grid formation. You can also select objects using the click-and-drag selection method, or via the **Edit—>Select All** or **Edit—>Select All Nodes** options on the primary window menu bar.
3. Once you've selected the objects you wish to align, release the **Ctrl** key and click on the Grid tool , or select **Layout—>Arrange—>Rectangle** from the primary window menu bar.

The selected objects will be aligned in a grid shape according to preset parameters, and placed in the upper-left-hand corner of the map window. Objects remain selected after the alignment is complete; you can move the grid as a whole by clicking and holding the mouse button, and dragging the grid of objects to the desired map location.

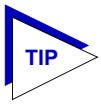

You can adjust the vertical and horizontal spacing of your grid using the Align and Space Evenly tools described in the previous sections.

Align Objects in Ring

Using the Ring tool , you can arrange selected objects into a single ring.

1. If necessary, open or activate the map which contains the objects you wish to align. (Remember, neither the Alignment Toolbar nor the Layout menu options will be active unless a map window is active.)
2. Press and hold the **Ctrl** key, and select the object(s) you wish to align in a circle formation. You can also select objects using the click-and-drag selection method, or via the **Edit—>Select All** or **Edit—>Select All Nodes** options on the primary window menu bar.

Note that all selected objects will be used to form the ring, which may be too large to be displayed in your map window. To prevent this from happening, use fewer objects in your ring, or adjust the spacing between objects, as described below.

3. Once you've selected the objects you wish to align, release the **Ctrl** key and click the Ring tool , or select **Layout—>Arrange—>Ring** from the primary window menu bar. The arrow cursor will turn into a crosshair cursor.

4. Place the crosshair cursor on the spot around which you want your ring to be centered, then click the mouse button to place the ring. The selected objects will be aligned according to the default spacing parameter, and a Ring Spacing window, Figure 6-30, opens.

Figure 6-30. The Ring Spacing Window

5. If the default spacing parameter does not produce the desired results, adjust the spacing by entering a new pixel value in the Ring Spacing window and clicking **Change**. The spacing will be adjusted, and the Ring Spacing window will remain available for additional changes. When you are satisfied with the appearance of your ring, click on **Quit** to close the Ring Spacing window.

Ring objects remain selected after the ring is formed; if you wish, you can click on any element in the ring and drag the ring to a new location. You can also move a ring using the right-mouse menu, as follows:

1. Make sure all objects in the ring (or all objects you want to include in the ring) are selected; if they're not, use the mouse to select them.
2. **Right-click** the mouse button on any of the selected elements, and select **Move\Create Ring** from the resulting menu. The arrow cursor will change to a crosshair cursor.
3. Move the crosshair cursor to the point on the map you wish to use as the center point of your ring(s), then click to place the ring.
4. The ring will be formed and the Ring Spacing window (Figure 6-30,) displays; adjust settings as required, then click **Quit** to close the window.

This feature can be handy if you have placed a ring in a location where its elements overlap other map objects or cannot be displayed in the map window; you can also move a ring formation by selecting all of its elements, clicking again on one element, and dragging the formation to a new location.

Using Object Linking and Embedding (OLE)

NetSight Element Manager can take advantage of Microsoft Windows' Object Linking and Embedding (OLE) architecture as a client application for data that has been created via external server applications. OLE is a convenient method of integrating information created via separate applications so that they can be accessed from a single location

(termed a “compound” document) — in this case, your network maps. You can use OLE to link or embed documents created via word processors, spreadsheets, databases, an audio application — or any other OLE server application — into your maps.

A linked or embedded data structure is generically termed an Object, since the data can be created by any OLE supporting application. From the NetSight Element Manager map interface, there is no apparent difference between a linked or embedded object — each appears on your map as an icon (or picture) within a border from which you can access the object data. The difference between an embedded or a linked object is in how its data is stored and updated with respect to the map data. Whether you choose to link or embed an object will generally depend on the circumstances in which people are going to access and use its information.

An object that is ***embedded*** accesses data that was created via an external application, but inserted *within* the map file data; the map becomes a container for the embedded data, which does not exist as an independent file. The benefit of embedding data is that when it is updated and saved, it is done within the confines of the map file data. For instance, you could embed voice annotations created via Microsoft Sound Recorder to describe each subnetwork map used by a single management station, and update them as necessary to reflect changes to your network. However, since the data is constrained to the map in which the object is embedded, it cannot be edited by its native application outside of the map file or shared (like linked data), and the size of the map file will increase correspondingly to the amount of embedded information.

An object that has been ***linked*** *references* data created from a source document that was created and saved external to the map. The object icon is a link or pointer to that source file. Clicking on the icon will open up the source document via the link, and any edits made to the data and saved will update the original. The advantage of linked information is that you can maintain a single source file and share its data via multiple links — for example, if you had multiple management stations running NetSight Element Manager, you could create and maintain a statistics database as a document on a central file server, then use a linked object in a map on each management station to reference that file. Each management station could copy statistics to that file that were applicable to their own subnetwork. However, since a source document is only maintained in one location, you must be careful that the links do not become corrupted by unintentionally deleting or renaming the source file. Objects are embedded by default; you must select a special option to create a linked object.

A separate means of linking and embedding information is to embed an Object Package into a map file. A “package” is a special type of embedded object — created via the Microsoft Object Packager — that itself contains an embedded or linked document, or part of that document, or an MS-DOS command line to be carried out from within the map file. You access the package contents by double-clicking on its icon, just as with other embedded or linked data. However, packaging data has some advantages over strictly linking or embedding data, since you can use the command line within a package to run a batch program or start another application from directly within a map.

NOTE

For more information regarding Object Linking and Embedding, as well as the Microsoft Object Packager, refer to the Microsoft Windows ***User's Guide for the Windows Graphical Environment****.*

Embedding an Object in a Map File

NOTES

This section provides directions for creating an embedded object icon ***from within*** *the map interface. You can also create an embedded object by working directly in the external source application, copy the data to the clipboard, and paste it into the map. For directions on how to work from the source application* ***before*** *embedding the object data in your maps, refer to your Microsoft Windows documentation, or to the documentation for your OLE-supporting application.*

For all directions on creating and saving a document in your source OLE application, refer to the appropriate documentation.

To embed an object within a map file:

1. Open or activate (as appropriate) the map file into which you want to insert an embedded object.
2. Select **Insert—>New Object** from the primary window menu bar. The Insert Object window (Figure 6-31) opens.

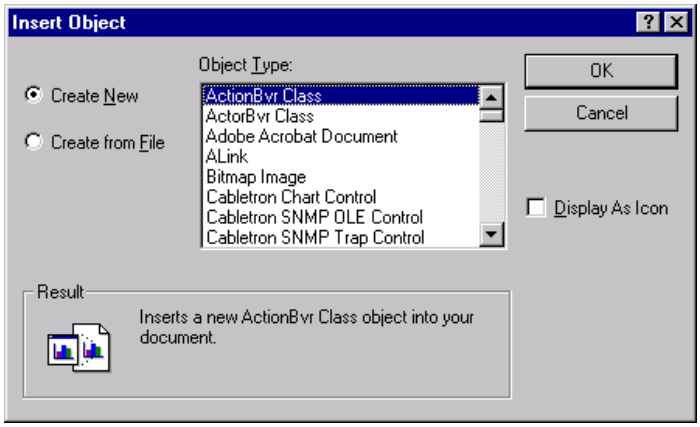

Figure 6-31. Sample Insert Object Window (Create New Selected)

3. The **Object Type** panel will contain all applications on your system that are registered with Microsoft Windows as server applications. Click to highlight the application Object Type of interest.

4. Objects on your map will be confined within a resizable border that contains a pictorial representation of the object data. The pictorial representation may be an icon to indicate the linked or embedded document, or the source data itself may be visible within the border. Resizing the border will scale the object's pictorial representation (the icon or the data).

   The pictorial display of an object will depend on its creating application type. For example, a Paintbrush picture object can display as the actual drawing or as a Paintbrush document icon; a Microsoft Word document or an object Package will always appear as an icon within the border.

   If you want to display the object as an icon within a border, or change the icon that will be associated with the embedded or linked object:

   a. Click to check the **Display As Icon** selection box in the Insert Object window (shown in Figure 6-33). The active icon associated with the object data type displays immediately below the field, along with a button allowing you to change the icon. By default, the object's active icon will be determined by the creating application.

   b. Click on **Change Icon**. The Change Icon window opens (as shown in Figure 6-32), with the active icon displayed in the lower right corner.

   c. Specify the icon and label that you want to use via the Change Icon window:

      1.) To leave the currently active icon in place, select the **Current** radio button.

      2.) To change the icon to the default icon associated with the object's creating application (if it is different from the current icon), select the **Default** radio button.

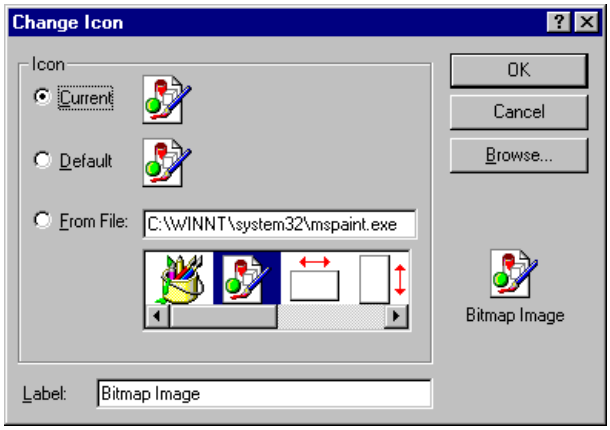

Figure 6-32. Sample Change Icon Window

3.) Select the **From File** radio button to change the icon to one that is associated with a different application or defined by a standard Windows icon file (with a .ico extension).

The icon(s) associated with the selected application file or defined by the icon file displays in the scrollable list box immediately below the file path text field. Click to highlight the icon of interest, or use the associated text field to type in the path and name of the icon you want to use. You can also click **Browse** and use the standard Browse window to search for and select a new icon file.

4.) In the **Label** text field, type in the descriptive text that will be displayed immediately below the icon in the map.

5.) Click **OK** to exit the Change Icon window and return to the Insert Object window.

5. **To embed** *new* **data** created via an external application into your map:

   a. Click to select **Create New**.

   b. Click **OK**. The external application will initialize, with a blank document window in which to enter new object data for embedding into the map.

   c. Enter in the data as desired.

   d. When you are finished entering the data, update the file to save the data and exit the application.

   e. A new embedded object displays within a border on your map.

You cannot create a link to an object created from within your NetSight Element Manager map file (as in the process described above); this data will always be embedded. If you wish to create a linked object, you must first create the object using its native application, save it as an independent file, then use the process described below and in the following section to link the object. Note that you must select the link option when you first insert the object, as you cannot select it once the object has been placed.

6. **To embed** *existing* **data that has previously been created** via an external application into your map:

   a. Click to select **Create from File**.

   As shown in Figure 6-33, a **File** field displays in which you can enter the path and name for the file whose data you want to embed; you can also use the **Browse** button to scan your directory tree for the file.

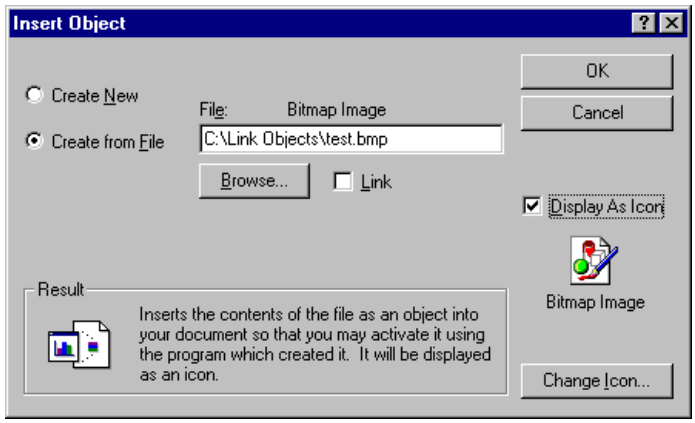

Figure 6-33. Sample Insert Object Window (Create From File Selected)

b. If you wish to link the existing object, be sure to select the **Link** option before placing the object in your map (see the following section for details), or click **OK** to embed the object in your map.

Creating an Object Link

If you wish to create a linked object, you must first create the object using its native application, save it as an independent file, then use the process described above to select the independent file as an object. Once you have selected the file you want to link, you must be sure to select the **Link** option in the Insert Object window (Figure 6-33) *before closing the window*. If you do not select the Link option when first placing the object, the object will be embedded, and you cannot change an embedded object to a linked object once it has been placed.

TIP

You can view information about your linked objects, change the source file for a linked object, and break the link to an object (making it a static map element) via the Links window; see ***Updating a Link****, page 6-54.*

Adding an Object Package

To add an Object Package to your map, make sure Package is selected in the Object Type panel, and then:

1. Click to select **Create New**.
2. Click **OK**. The Microsoft Windows Packager application will initialize, with a window in which to create a new package to be embedded into the map. The title bar of the Package window will read “Object Packager: Package in XXX Map” (see Figure 6-34).

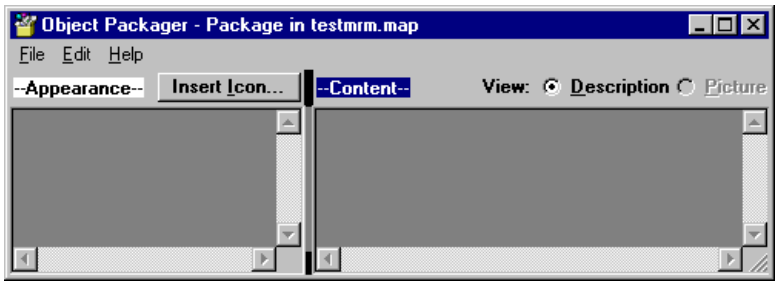

Figure 6-34. Sample Object Packager Window

3. Create an Object Package. Embed or link an entire document (or part of a document) into the package, or build an MS-DOS command line to run a batch file or start an application via the package; then specify the icon and label to be used for the package. The details of how to create an Object Package are provided in your ***Microsoft Windows User's Guide***.
4. Select **File—>Update** from the Object Packager window to save the Object Package, and then select **File—>Exit** to close the Object Packager window.
5. A new embedded Object Package displays in your map, using the specified icon and label.

Editing Object Data

To open the map object's data in the source application that created it (or to launch the contents of a packaged object) so that you can edit or update the object data if necessary:

1. Double-click within the object's border.

   or

   Single click to select the object, then choose **Edit—>Package Object—>Activate Contents** from the primary window menu bar.

NOTES

If an object has been opened from the map so that the source application is active, you will notice that the object will be shaded in the map to indicate its active or opened state.

Before you update an object's data, remember:

- *If the object is **embedded**, updates to its data are written directly within the map file.*
- *If the object is **linked**, updates of its data will be written to the source document and reflected in **all** links that access the source document.*
- *If the object is **Packaged**, the form in which the data is updated will depend on whether it was linked or embedded within the Package.*

Updating a Link

You can use the Links window to see a list of all linked objects in any map, and:

- Specify whether to manually or automatically update the appearance of the linked object data within its border (if it is not displayed as an icon).
- Open the source document of any linked object.
- Change the linked object to reference a new source document.
- Remove the link to the source document, so that the object becomes a static item in your map.

To access the Links window from any open map:

1. Make sure the desired map is active, and then select **Edit—>Links** from the primary window menu bar. The Links window (Figure 6-35) opens.

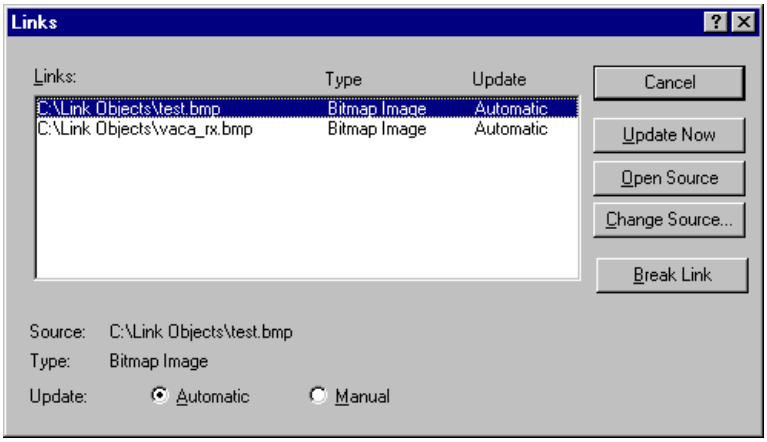

Figure 6-35. Sample Links Window

The Links window will display each link object in the active map. The specifics of each link will be shown in a panel in the window. There are also link-specific fields at the bottom of the window, and an array of edit buttons (**Update Now**, **Open Source**, **Change Source**, and **Break Link**) to the right of the window.

The fields at the bottom of the window and the edit buttons will be inactive until you select a link to edit. To select a link so that you can make changes to it via the fields and buttons, click to highlight it in the panel.

Links are identified by the following fields and panel columns:

Links — This column identifies the file path and name for the source document referenced by the link; if the link is selected, the file path and name will also display in the **Source** field.

Type	This column indicates the source document's type; if a link is selected this will also display in the **Type** field.
Update	Indicates whether the appearance of the object in the map will be automatically or manually updated to reflect changes to the source document; if a link is selected, this will also be indicated by the **Automatic** or **Manual Update** radio buttons at the bottom of the window.

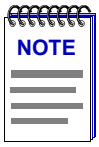

Remember, the Update state of a link determines how it appears in the map — and does not correspond to the version of the source document that is accessed via the linked object. Opening the source document through the link will always open its last saved version.

To determine whether the object's appearance will be updated automatically or manually:

1. Click to highlight the desired link object in the panel.
2. Click on the appropriate radio button option at the bottom of the window: **Automatic** or **Manual**:

   If you select **Automatic**, the appearance of the link object in the map will be updated each time the source document is updated and saved.

   If you select **Manual**, you determine when to update the appearance of the link by clicking on **Update Now**. Each time this button is clicked, the object's appearance will be revised to reflect the latest saved version of the source document.

To open the source document for a link, so that you can edit its data:

1. Click to highlight the desired link object in the panel.
2. Click on **Open Source**. The source document will open in the application that created it. When you close the source application window, you will be returned directly to the map window.

To change a link's reference to a source document, so that opening the object will link to a newly specified file path or name (in case the source document was moved or renamed, or if you simply want to reference a different document via the link):

1. Click to highlight the desired link object in the panel.
2. Click on **Change Source**. A Change Source window displays, in which you can browse your directory tree to select a new file reference.

Enter a file path and name directly in the **Source** text field at the top of the window, or use the **Directories** tree and **Drives** drop-down list box to scan your system directories; and use the **List Files of Type** drop-down list box and files panel to select a source file. Click to select a file and enter it in the Source text field.

You can also remove a link between the object and a source document, so that the object no longer opens data via an external application. The object will still be pictorially represented in your map as it was prior to breaking the link, and can still be resized or moved around on the map; however, attempting to open the object to access data will result in a message stating "Warning: Cannot activate a static ActiveX object."

To break a link:

1. Click to highlight the desired link object in the panel.
2. Click on **Break Link**. The object will immediately lose its reference to a source document, and the link will be removed from the Links window.

Converting an Object

More recent Microsoft Windows applications (those that support the OLE2 version) may allow you to convert an embedded object from one similar application type to another. This is useful if you intend to be sharing data between machines in one form or another. For example, you might copy a map with an embedded spreadsheet from one machine, and then install a copy of that map onto a second machine that uses a different spreadsheet application than the original machine. By using conversion, the second machine could open the embedded object in its native spreadsheet format.

There are two options for converting an embedded object:

- You can permanently convert object data from its original source format to that of another supporting application. When the conversion is complete, the object will open into the new document format; the original document format will be overwritten and will no longer be available.
- You can choose to emulate the object data via a new supporting application. Emulation means that the data will be temporarily treated as if it had been converted, but when the source file is closed it will retain the data structure of the original creating application. This is useful if, for example, you have several management stations that run different applications, but share a common link to a source document which you want to maintain in its original format.

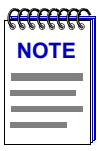

You cannot convert a linked object via NetSight Element Manager; linked objects must be converted from within their source application.

To convert an embedded object to a new application format:

1. Open or activate the map which contains the object you wish to convert, and click to select the object.
2. From the primary window menu bar, select **Edit—>XXX Object—>Convert**. The Convert window, Figure 6-36, opens.

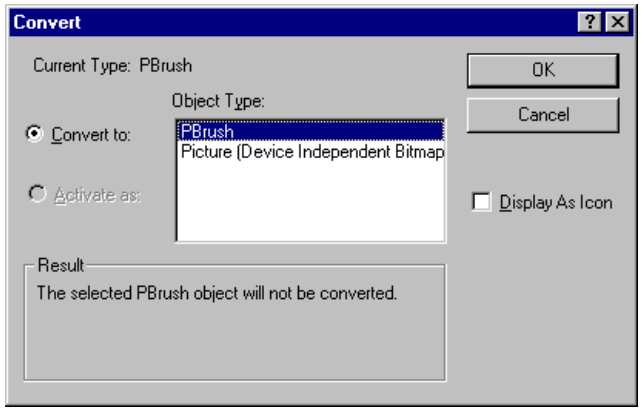

Figure 6-36. Sample Convert Window

3. In the **Object Type** panel, click to highlight the application type to which you want to temporarily or permanently convert the embedded or linked data.
4. Select the **Convert to** radio button to specify that the object's source document be transformed into the new application type.

   or

   Select **Activate as** to specify that the source document be opened in the new application type, but retain its original format when it is updated and closed.

   The Result message text will reflect your selection.
5. Click **OK** to begin the conversion process. The Convert window will close, and if a permanent conversion was specified, the old data will be overwritten into the new file format.

The next time you access the object data, it will be opened in the specified application type. You can also use the Convert window to change the icon to be used to represent the object in the map, as detailed on page 6-49.

Working with Linked and Embedded Objects

You can manipulate data objects in your map just like any other map object. They can be resized, annotated with notes, layered by moving forward or backward, and so forth.

You can also perform the standard **Cut**, **Copy**, and **Paste** options on embedded or linked objects in your map. If you copy an embedded object and paste it to another map, the source data will be copied into the new map; if you copy a linked object, the source file pointer will be copied into the new map.

Cutting or deleting an embedded object (including an Object Package) will remove its pictorial representation from the map and strip out the application data; cutting or deleting a linked object will remove the pointer to the source document (although the document itself will remain at the location from which it was linked).

Editing Your Map

You can use the standard **Cut**, **Copy**, **Paste**, **Delete**, and **Select All** options on the primary window **Edit** menu, the standard Toolbar, and the map right-mouse menu to perform basic map editing functions; you can also use the Delete key on your keyboard, and the *ctrl*-drag-and-drop function for copying (hold down the Ctrl key while dragging and dropping to copy rather than move an object). These functions perform the way you would expect them to for any object — icon, symbol, text, link, or graphic element — in your map. However, there are a couple of special cases you should be aware of.

First, note that deleting a device from a map does ***not*** necessarily delete that device from the central node database; when you delete the last instance of a device icon that appears in any open map, you will be offered a series of choices as to how far you want the delete to go (see Figure 6-37): deleting from the **Active View** (the default selection) removes the device from the selected map only; deleting from **All Views** moves the node into the Excluded Nodes database; and deleting from the **Database** removes the node entirely. See Chapter 4, **List Views**, for details.

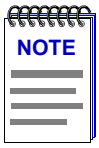

*If you select the **All Views** or **Database** option for a device which is still represented on a map which is closed, that device's icon will be in an unknown state the next time you open that map. You can either delete the icon or re-define it using the appropriate IP address; note that re-defining the icon will re-add the device to the central node database.*

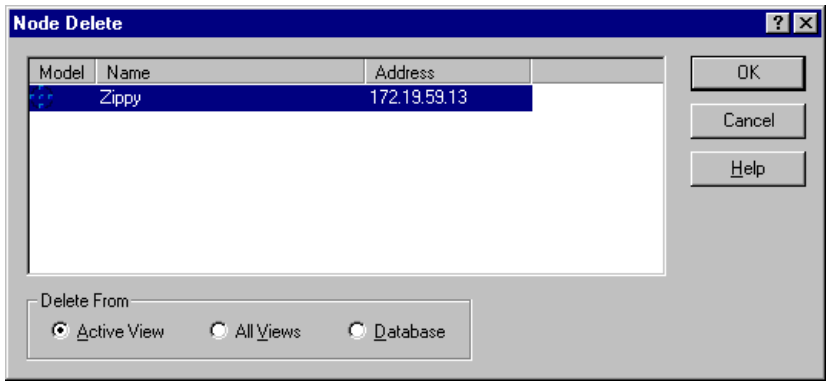

Figure 6-37. The Node Delete Window

Second, you can cut and paste or move Submap icons, and the links will change accordingly; remember, however, that no map can be a submap to more than one root map. If you try to place multiple submap icons representing the same map, the submap link will change without warning (and the workspace display will update accordingly). See **About Submap Icons**, page 6-3, and **About GoTo Symbols**, page 6-4, for more information.

You should use care when removing a submap link entirely. If you remove a submap icon with a **Cut** command, only the icon (and the link it represents) will be removed; the actual map file remains intact, and can be pasted into another map as a **Go To** link, if you wish. However, if you use a **Delete** command (or the Delete key on your keyboard) to remove a submap icon, a prompt window opens, asking if you wish to remove the view. If you select **Yes**, the .MAP file represented by the submap icon will be deleted, as will the submap icon itself. If you select **No**, only the submap icon (and the link) will be removed.

TIP

*If you inadvertently remove a .MAP file along with its submap icon by selecting **Yes** in the prompt window, don't worry. The associated .MRM file will still be in place, and can be re-imported. See **Importing an MRM Map**, page 6-20, for more information.*

Third, when a map window is the current active window, the **Edit** menu contains an additional, non-standard selection: **Paste Special (No IPs)**. You can use this option to copy and paste single icons or a group of icons without including each icon's descriptive information. This feature allows you to repeat a pattern of device icons on multiple maps, and still allows you the freedom of assigning the appropriate IP address (and any other descriptive information you need) for each subsequent map.

To use the Paste Special option:

1. Create, open, or activate the appropriate map window(s), if necessary.
2. Use the **left** mouse button and the **Ctrl** key (or click and drag) to select the icon or icons whose arrangement you wish to repeat.

3. Select **Edit —>Copy** from the primary window menu bar, the **Copy** button on the toolbar, or the **Copy** option on the right-mouse map menu.
4. Create, open, or activate the map to which you wish to add the copied icons, and select **Edit —>Paste Special (No IPs)** from the primary window menu bar.
5. Place the **No IP Paste** cursor in the location to which you wish to paste the icon group, and click the mouse button to paste. (Note that the bounding box surrounding the cursor indicates the size of the icon group you have copied.) The icons display on the map in their original arrangement, but will be labeled N/C (not connected).
6. To define each icon, double-click the undefined symbol; the standard Insert Device window (Figure 6-17, page 6-24) opens. Proceed as described in **Adding a Device Directly to a Map**, beginning on page 6-23. If necessary, the icon will update automatically to reflect the appropriate node type.

Fourth, when you use the cut and paste functions to move a node icon or any graphic element in a map, clicking **Paste** will change the cursor to a special paste cursor; you can move this paste cursor around on the map or onto another map, then click again to paste the element.

And finally, note that the **Edit** menu contains two selection options: **Select All** will highlight every element in the active map; **Select All Nodes** will highlight only the device (or "node") icons.

Protecting Your Maps from Editing: The Lock Feature

Once you have configured your maps with all the node icons and other graphic elements you need, you can protect them from accidental editing by using the **Lock** feature. Locking a map allows only three actions to be performed: you can single-click an icon to select it (as a means of launching any applicable tool against it), right-mouse click an icon to access the **Manage** option on the resulting menu, or you can double-click it to launch the Chassis Manager application.

To lock a map:

1. With the map you wish to lock open and active, select **View—>Lock** from the primary window menu bar.

   or

   Select **File—>Open Map** from the primary window menu bar. In the resulting Open window, select the **Open as read-only** option. The map will open in a locked state.

   or

   Open the Map View Properties window for the map of interest, click to select the **Lock View** option, then click **OK** to close the Properties window and SET your changes.

When a locked map is the active map, the map icon in the upper-left-hand corner of the map window will change into a lock; also, a lock icon will precede the map name in the workspace, and the **Lock** option on the **View** menu will be preceded by a check mark. To unlock the map, click to select the **Lock** option again.

A map's locked state will persist until it is changed by user action.

Managing Maps and Folders

The following sections describe how to perform general housekeeping functions — opening, closing, saving, and deleting — for maps and folders.

Opening and Closing Maps and Folders

There are several ways to open an existing map:

1. Select **File —>Open Map** from the primary window menu bar.

   or

   Click on the toolbar.

2. The standard Windows File Open window displays; select the appropriate directory (if necessary), then double-click the appropriate file name (or click to select the file name, then click **OK**).

You can also open a map by double-clicking the appropriate Submap or GoTo symbol, or by double-clicking the map name in the Map View workspace.

Similarly, there are a number of ways to close a map window:

1. Select **File —> Close Map** from the primary window menu bar.

   or

   Click on the window close control located in the upper-right corner of each map window.

No matter what closing method you choose, you will be prompted to save the map if it has changed since it was opened. You will also be prompted to save all open maps any time you exit NetSight Element Manager; see **Saving Maps**, below, for some important information about saving maps.

Since folders are a Map View display convention only, all controls related to folders are accessed from within the Map View workspace. To open or close a folder, click once on the plus (+) or minus (-) sign located to the left of the folder name, or double-click on the folder icon.

Saving Maps

When NetSight Element Manager saves a map file, it actually saves it in two distinct formats: a .MAP format, which is the format in which you open and edit maps within the primary window; and a .MRM format — a text-based format which allows you to recover corrupted or accidentally-deleted .MAP files, or allows you to move a map from one folder to another (see **Importing an MRM Map**, page 6-20, for more information.) Both versions of each map file are stored by default in the directory you defined via the **Tools—>Options—>Directories** settings; however, the .MRM files will not be listed when you select an Open Map option, as the .MAP files are the ones viewed in NetSight Element Manager.

To save a map file:

1. Click on the Save Map button on the toolbar.

   or

   Select **File —>Save Map** or **File —>Save Map As** from the primary window menu bar.

Note that, with either of these procedures, **only the current active map will be saved!** To ensure that all changes to all maps are saved, you must activate each map in turn and save it via one of the options above, or use the Save All Maps function:

1. Click on the Save All Maps button on the toolbar.

   or

   Select **File —>Save All Maps** from the primary window menu bar.

As a safety precaution, you will always be prompted if you attempt to close a map or exit NetSight Element Manager without saving your changes.

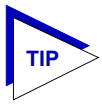

*If you attempt to save a locked map, you will be offered the option of saving a **copy** of the locked map. You must assign this copy a unique name, or the save will not take place.*

Deleting Maps and Folders

You can delete both maps and folders via the Map View workspace, with three restrictions: you cannot delete a locked map; you cannot delete any map which is serving as the root map for one or more submaps; and you cannot delete a folder which contains one or more maps.

1. In the Map View workspace, click the **right** mouse button on the map or folder you wish to delete.
2. Select **Delete** from the resulting menu.

The selected map or folder will be deleted without further confirmation. Note that deleting a map deletes only the .MAP version of the map file; the .MRM version remains intact in the assigned directory and can be re-imported at any time. (To remove lingering .MRM files, use the Windows Explorer or any other Windows file management application.)

Chapter 7

Node Management Overview

Device management overview; using Source Address Naming; using the PING tool

This chapter provides a brief overview of some of the management capabilities provided by NetSight Element Manager, including device management, source address naming, and the use of the PING tool.

Device Management

NetSight Element Manager provides extensive device management capabilities through its Device Manager and Chassis Manager applications. This management capability is described in detail in the device-specific ***User's Guides*** included with your documentation; this section provides some brief introductory information.

Launching Device Management

The primary interface NetSight Element Manager provides for device management is the Device View (also called the Chassis View). This view serves as a single point of access to all other device-specific management functions, and provides a graphical view of the device you are managing, the hub in which it is contained, and the status of its individual interfaces. There are several ways you can launch a Device View window:

1. In any map, list, or tree view, double-click on the icon representing the device you wish to manage.

   or

1. In any map, list, or tree view, select the device you wish to manage.
2. Select **Manage—>Node** from the primary window menu bar, or select the Manage Node toolbar button.

   or

1. In any map, list, or tree view, click the **right** mouse button once to select the device you wish to manage.
2. On the resulting menu, click to select **Manage**.

You can also access Device Manager and Chassis Manager from the Start menu:

1. From the Start menu, select **Programs > Enterasys Networks > NetSight Element Manager > Device Manager** (or **Chassis Manager**).
2. Enter the desired IP address and community name, and click **OK**. The Device View opens.

Once a Device View window is launched, it will operate as described in the device-specific ***User's Guides***.

A Brief Overview of Device Management

After you launch device management, the appropriate Chassis View or Device View window (as illustrated in Figure 7-1) opens. This window is the color-coded graphical interface that provides you with immediate feedback on device configuration and status, and lets you access further device management windows.

For devices which NetSight Element Manager specifically supports, a description of the device, along with its IP address, displays in the Module or Chassis View window title. Menus at the top of the front panel display, or available from the interface display itself, allow you to manage the device at various physical levels (device, board, or port) and functions (FDDI, repeater, or bridge), and provide access to a variety of utilities.

If NetSight Element Manager can only communicate with a device via standard MIB I or MIB II objects, it will have no device-specific management available. In this case, a "generic" window opens.

Figure 7-1 shows three management windows: both a Module View and a Chassis View window for the SmartSwitch, and a "generic" Chassis View window for use with an SNMP device which is not specifically supported by NetSight Element Manager.

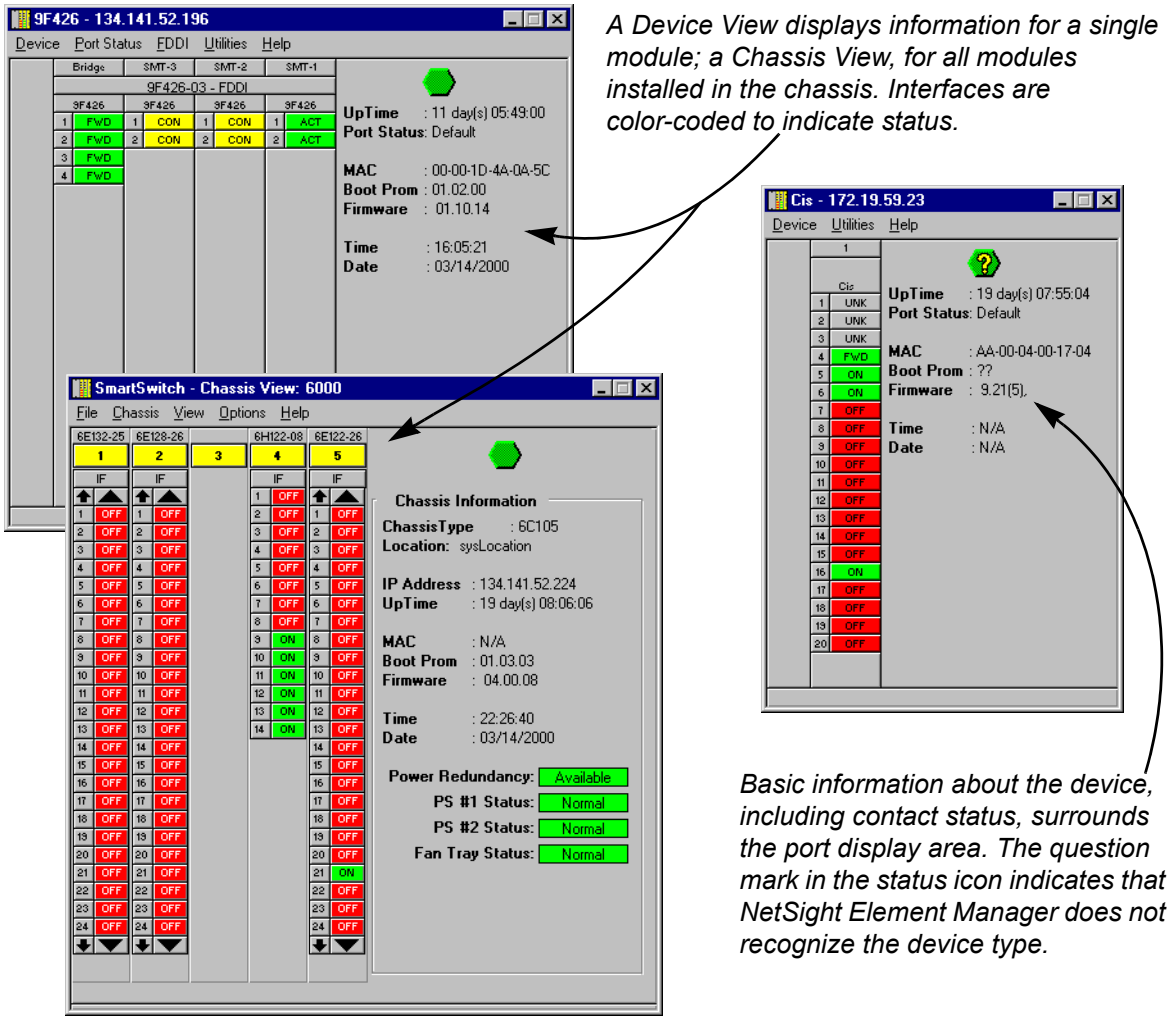
A Device View displays information for a single module; a Chassis View, for all modules installed in the chassis. Interfaces are color-coded to indicate status.
Basic information about the device, including contact status, surrounds the port display area. The question mark in the status icon indicates that NetSight Element Manager does not recognize the device type.
9F426 - 134.141.52.196
Device Port Status FDDI Utilities Help
UpTime : 11 day(s) 05:49:00
Port Status: Default
MAC : 00-00-1D-4A-0A-5C
Boot Prom : 01.02.00
Firmware : 01.10.14
Time : 16:05:21
Date : 03/14/2000
Cis - 172.19.59.23
Device Utilities Help
UpTime : 19 day(s) 07:55:04
Port Status: Default
MAC : AA-00-04-00-17-04
Boot Prom : ??
Firmware : 9.21(5),
Time : N/A
Date : N/A
SmartSwitch - Chassis View: 6000
File Chassis View Options Help
Chassis Information
ChassisType : 6C105
Location: sysLocation
IP Address : 134.141.52.224
UpTime : 19 day(s) 08:06:06
MAC : N/A
Boot Prom : 01.03.03
Firmware : 04.00.08
Time : 22:26:40
Date : 03/14/2000
Power Redundancy: Available
PS #1 Status: Normal
PS #2 Status: Normal
Fan Tray Status: Normal

Figure 7-1. Sample Module and Chassis View Windows

Any network interfaces that NetSight Element Manager detects on such a device will be color-coded for SNMP MIB-II *ifOperStatus* (interface operational status): Up (green), Down (blue) or Test (magenta). Any identifying fields provided for the device (e.g., an interface description available from the Port menu, or the Device Type available from the Device menu) are taken directly from SNMP MIB-II information stored at the device, and management is limited to the Generic SNMP application. Refer to your ***Generic SNMP User's Guide*** for details.

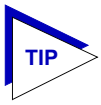

The Chassis Manager application will generate an event when it is used to perform an SNMP SET; this event will record the MIB object and instance that was changed for tracking administrative actions on devices. All SNMP SETs made through Chassis Manager will be recorded in the Alarm and Events Manager's cache by default. You can turn off logging of device SETs by using a text editor to edit the CTRON.INI file (located in the winnt or windows directory); simply look for the line ***LogSets=on*** *and change it to* ***LogSets=off****, or erase the line completely. You can re-start this logging by re-adding the line (or re-setting its value); be sure to enter the correct case when retyping or editing the line.*

Source Address Naming

Most devices support a variety of source addressing features, all based on the ability of each device to create and maintain a database of the MAC addresses which are communicating through each of its ports. NetSight Element Manager allows you to view and control these databases in a variety of ways, including:

- Examining a list of source addresses (or user names) that are being processed by any port on the selected device.
- Maintaining security on a device by locking ports from source addresses that are not currently contained in their database.
- Enabling you to find the port through which a source address is communicating.
- Defining an ageing time after which a non-communicating source address will be removed from the source address database.

Each of these features is discussed more thoroughly in the individual device-specific ***User's Guides***; in addition, however, NetSight Element Manager also provides a global source address function called Source Address Naming. With source address naming, you can create a text file which will add names to correspond to the source addresses that appear in any device- or port-specific Source Addresses window.

To add source address names to the addresses appearing in the Source Addresses window:

1. Create a text file named ENETADDR.DEF in NetSight Element Manager's data directory: **\NetSight Element Manager x.x\bin**.

   The ENETADDR.DEF file should have the following format for each listed device:

   <MAC address> <logical name>; optional comment

   for example:

   00001d027f18 Bill's PC; 486-66 ACME Super PC

NOTE

For this function to work your ENETADDR.DEF file must be saved in a text-only format.

Figure 7-2 displays a sample ENETADDR.DEF file.

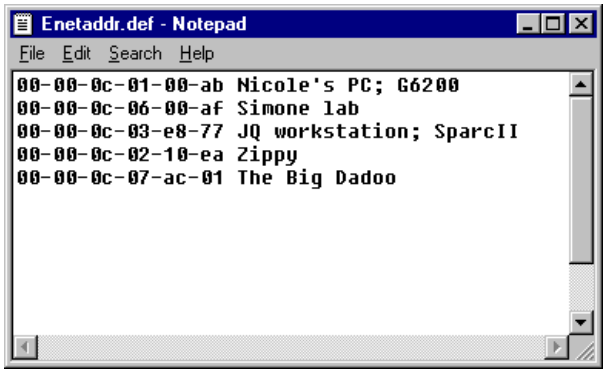

Figure 7-2. Sample ENETADDR.DEF File

2. Edit the **ctron.ini** file by adding the following line under the heading [SpectrumElementManager2]:

   SrcAddrFileName=<pathname><ENETADDR.DEF>

   where <pathname> is the path to the data directory in NetSight Element Manager (by default, \NetSight Element Manager x.x\bin).

3. Save and close the **ctron.ini** file

4. Exit and re-start NetSight Element Manager, so that the changes to **ctron.ini** will take effect.

NOTE

*You need not close, then re-open NetSight Element Manager after making edits to the ENETADDR.DEF file; it's only necessary after editing the **ctron.ini** file.*

Once you have configured this file, any Source Address window (like the one illustrated in Figure 7-3) will include the text name you have assigned to a MAC address, as well as the MAC address itself.

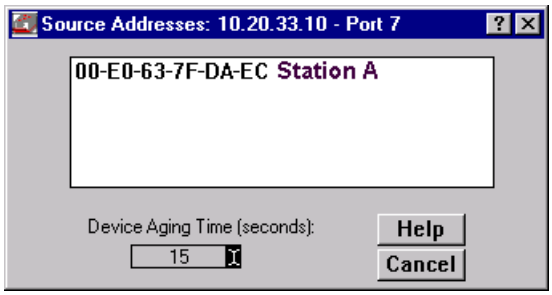

Figure 7-3. Sample Source Addresses Window

TIP

If you have Network Analyzer software, you can use the ENETADDR.DEF file generated by the Analyzer for all source addresses. Simply copy the existing file into the NetSight Element Manager data directory (by default, \NetSight Element Manager x.x\bin).

If you are upgrading from a previous version, your existing ENETADDR.DEF file will automatically be saved and restored to the bin directory.

PING

A PING (Packet INternet Groper) is an ICMP (Internet Control Message Protocol) echo request, in which one node sends the message request "are you alive?" to another node, and awaits a reply. If a message is returned, the receiving device has been successfully "PINGed." NetSight Element Manager's PING tool allows you to use this quick and easy test of device status and/or network connectivity.

Accessing the PING Window

There are two ways to access the PING tool in NetSight Element Manager:

1. Select the node you want to contact via PING in a list, tree, or map view.
2. Choose **Tools —> Ping** from the primary window menu bar.

   or

   Click the PING icon in the toolbar.

   The PING window, Figure 7-4, opens.

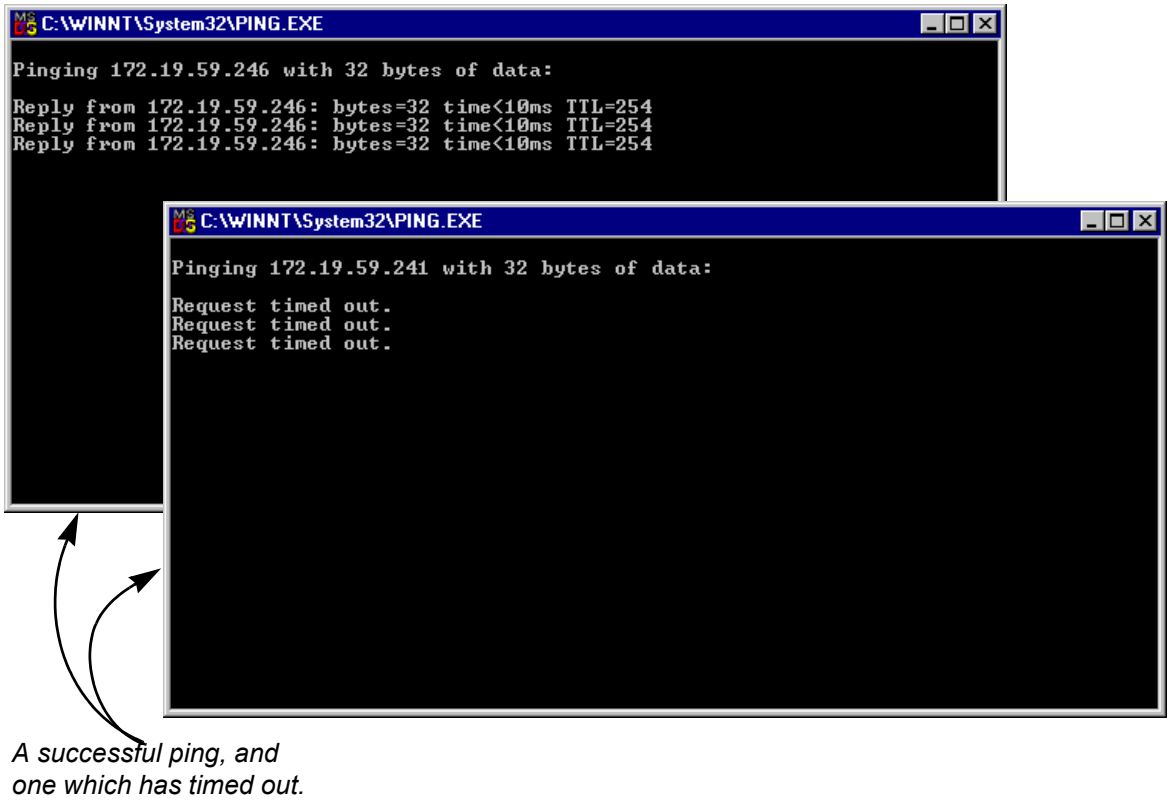
C:\WINNT\System32\PING.EXE
Pinging 172.19.59.246 with 32 bytes of data:
Reply from 172.19.59.246: bytes=32 time<10ms TTL=254
Reply from 172.19.59.246: bytes=32 time<10ms TTL=254
Reply from 172.19.59.246: bytes=32 time<10ms TTL=254
C:\WINNT\System32\PING.EXE
Pinging 172.19.59.241 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
A successful ping, and
one which has timed out.

Figure 7-4. The PING Window

NetSight Element Manager sends a PING request to the selected device, and the response, if any, is displayed in the PING window.

Appendix A

Batch Customization via AddTool, AddImage, and AddDevType

Using the AddTool, AddImage, and AddDevType tools to expand the capabilities of NetSight Element Manager

Three options — the ability to add an option to the Tools menu, the ability to define a new Node Class to represent a device type, and the ability to define (and add an icon representing) a new Model, Class, Enterprise, or Topology value — are available external to the NetSight Element Manager framework via three executables: AddTool, AddDevType, and AddImage. These programs process input from standard text files and allow for quick batch customization of these functions. The use of each tool is described in detail in the following sections.

Using the AddTool Utility

The AddTool utility allows you to add Tools menu options and Toolbar icons to the NetSight Element Manager framework, as described in **Customizing the Toolbar** in Chapter 2, **Overview of NetSight Element Manager**. To do so, you must prepare an input file in one of two formats — comma-separated variable (CSV), or tag/data format — which contains the following information for each tool you wish to add:

Menu Text

This parameter defines the text that will appear in the Tools menu; you can add an ampersand (&) in front of the character you wish to use as the keyboard mnemonic for your new tool option. Note that you cannot duplicate any mnemonics already in use on the Tools menu.

Hint Text

This parameter defines the text description that will appear in the primary window status line when the mouse pointer is held over the associated Toolbar icon or Tools menu pick.

Bubble Text

Bubble text appears in a pop-up window when the mouse pointer is held over the associated Toolbar button.

Command

This parameter specifies the executable that will be launched when the associated Toolbar button or Tools menu item is selected. Note that you must specify a file which exists, or the AddTool operation will fail; you should also specify the full path to the executable.

Arguments

You can use this optional parameter to specify any arguments or options that should be passed to the specified command. You should specify any necessary switches in the format required by the application you want to add; if required or allowed, you can also add a specific argument string, or one or more of several available substitutions:

%I	inserts the IP address of the selected node
%C	inserts Community Name associated with the selected node

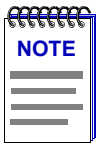

Note that if you use the IP address or community name argument substitutions, the associated Toolbar icon and Tools menu option will only be activated when a node is selected in a list, tree, or map view.

%O	if a node is selected, inserts its *sysOID*
%N	if a node is selected, inserts its assigned name. Note that, for a device, this will be either the *sysName* value set at the device, or the Name value set via the node's Properties window.
%H	inserts the parent window identifier.
%%	inserts a single percent sign

Working Directory

Use this optional parameter to specify the location of any secondary files that might be needed by the executable, if those files are not located in the same directory as the executable itself. If you choose not to specify a working directory, the directory in which the executable resides will be used by default.

Small Image File

You can use this optional field to specify the small (16x16 pixel) bitmap image file that will be used to represent your tool on the Toolbar icon. Be sure to specify the path to the bitmap image. If you choose not to specify an image file, the default icon associated with the application will be used.

Large Image File

You can use this optional field to specify the large (32x32 pixel) bitmap image file that will be associated with your application. Be sure to specify the path to the bitmap image. If you choose not to specify an image file, the default image associated with the application will be used. Note that the large image file is not displayed anywhere in the NetSight Element Manager framework.

Creating an Add Tool Input File

You can create your Add Tool input file in one of two formats: multi-line tag/data format, or single-line comma-separated variable (CSV) format. Each method is described below.

Comma-Separated Variable (CSV) Format

To create an input file using the single-line CSV format, you need only string together the fields you wish to use in a single line, with each value separated by a comma. Spaces before and after commas are ignored. If you wish to skip one or more optional fields, be sure to use a comma for a placeholder, or any following fields will be misread. Use one line for each tool you wish to add; extra lines are ignored. Note that the filename cannot contain spaces.

If you wish, the first line of your file can contain a comma-separated list of the available field names; if you choose not to include this line, you must launch your CSV file with a special switch (as described in **Executing the AddTool Command**, page A-4). A CSV file *with* this descriptive line is the default file type expected by the AddTool utility.

The following CSV file would add the MIB Tools and Notepad utilities to the Tools menu and toolbar. This sample includes the optional initial descriptive line, and demonstrates (in the Notepad line) using commas as placeholders for the optional fields you do not wish to use. Note that the line breaks shown here are a publishing limitation only; be sure your file does not break the single line required for each tool entry.

CSV Sample File

```
Menu text, Hint text, Bubble text, Command, Arguments,
Working directory, Small image file, Large image file

MIB &Tools, Launch the MIB Tools application, MIB Tools,
MIBMGR.EXE, /ip %I /comm %C, D:\NetSight Element Manager
x.x\DATA, smlLeaf.bmp, lrgLeaf.bmp

&Notepad, Launch Notepad, Notepad, C:\WINNT\notepad.exe, , ,
C:\images\small_note.bmp, C:\images\large_note.bmp
```

Tag/Data Format

To create a file using the multi-line tag/data format, you must place each field name and the value you want to use in its own line. Field names and their values are separated by a colon with a space on either side. Because each line of data is tagged with its function, you

need not use placeholders for optional fields you do not wish to use. The end of each entry is signaled by a separator line containing three dots; any additional blank lines between entries are ignored. Note that the filename cannot contain spaces.

If you choose to use the tag/data format for your input files, you must launch the AddTool utility with a special switch, as described in **Executing the AddTool Command**, page A-4.

The following tag/data formatted file would add the MIB Tools and Notepad utilities to the Tools menu and toolbar; it also demonstrates that optional fields you do not wish to use need not appear in the tool entry.

Tag/Data Sample File

```
Menu text : MIB &Tools
Hint text : Launch the MIB Tools application
Bubble text : MIB Tools
Command : MIBMGR.EXE
Arguments : /ip %I /comm %C
Working directory : D:\NetSight Element Manager x.x\DATA
Small image file : smlLeaf.bmp
Large image file : lrgLeaf.bmp
...
Menu text : &Notepad
Hint text : Launch Notepad
Bubble text : Notepad
Command : C:\WINNT\notepad.exe
Small image file : C:\images\small_note.bmp
Large image file : C:\images\large_note.bmp
...
```

Executing the AddTool Command

Once you have created the necessary input files, you can run the AddTool utility via the Run window, as follows:

1. Make sure NetSight Element Manager and all of its components are shut down; the AddTool utility will not run if NetSight Element Manager is running.
2. Select **Start—>Run** from the Windows **Start** button, then enter the following command in the **Open** field:

   ```
   C:\NetSight Element Manager x.x\bin\AddTool.exe [/M] [/D]
   input_filename
   ```

   where:

   /M or **/m** signals that the input file is in multi-line tag/data format

   /D or **/d** signals that the input file is in CSV format, but does not contain the initial descriptive line

   input_filename specifies the path and filename of your input file

3. Click **OK** to launch the utility.

If the tool did not run successfully, an error message will appear, indicating the source of the problem. Note that the tool stops executing at the first error, so if your file contains multiple errors, they will only be detected one at a time, each time you run the tool If no message windows appear, the tools have been added successfully.

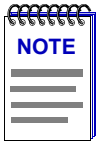

*Once you have added tools via the AddTool utility, they can only be removed via the **Tools—>Customize** window available from within the framework. See **Customizing the Toolbar** in Chapter 2, **Overview of NetSight Element Manager**.*

Using the AddImage Utility

NetSight Element Manager classifies devices according to what is called a Node Class. Each defined Node Class specifies four kinds of attributes that describe the device: **Model** (the model type of the device, such as CSX200 or 6H123-50), **Class** (the functional classes supported by the device, such as Switch, Repeater, etc.), **Topology** (topologies supported by the device, such as FDDI, Token Ring, or Ethernet CSMACD), and **Enterprise** (the manufacturer of the device). NetSight Element Manager provides many pre-defined model, class, topology, and enterprise definitions, but if you wish you can use the AddImage tool to add any additional ones you might need. The AddImage tool performs the same function as creating a new image via the Node Class Properties window (as described in **Editing Node Class Properties** in Chapter 4, **List Views**).

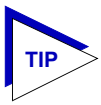

*For more information about creating an entirely new node class, see **Using the AddDevType Utility**, page A-9.*

As with AddTool, to use the AddImage utility you must prepare an input file in one of two formats — comma-separated variable (CSV), or tag/data format — which contains the following information for each image you wish to add:

Image Type

Use this parameter to specify which type of image you wish to add: Model, Class, Topology, or Enterprise. Each image can be assigned only a single type, but you can create a file which contains definitions for more than one type.

Image Name

Use this parameter to specify the name that will be assigned to your image. Note that the case of this name is significant; an image named “superzippy” will not be found in a device type definition (see **Using the AddDevType Utility**, page A-9) that references an image named “SuperZippy.”

Small Image File

Use this field to specify the small (16x16 pixel) bitmap image file that will be used to represent your image in a variety of places within the NetSight Element Manager framework (in the appropriate Select Symbol window, in a List View or Tree View, etc.). Be sure to specify the path to the bitmap image.

Large Image File

Use this field to specify the large (32x32 pixel) bitmap image file that will be used to represent your image in a variety of places within the NetSight Element Manager framework (in the appropriate Select Symbol window, in a List View or Tree View, etc.). Be sure to specify the path to the bitmap image.

Image Identifier

Use this field to assign a unique numeric identifier for your image. This identifier must be unique among the identifiers already defined for the image type (model, class, topology, or enterprise) you are defining. You can check to see which identifiers have already been used for each image type via the Node Class Properties window; see **Editing Node Class Properties** in Chapter 4, **List Views**, for more information.

Creating an AddImage Input File

You can create your AddImage input file in one of two formats: multi-line tag/data format, or single-line comma-separated variable (CSV) format. Each method is described below.

Comma-Separated Variable (CSV) Format

To create an input file using the single-line CSV format, you need only string together the fields you wish to use in a single line, with each value separated by a comma. There are no optional parameters, so be sure you have included the correct number of values. Spaces before and after commas are ignored. Use one line for each image you wish to add; extra lines are ignored. You can include different image type definitions (model, class, etc.) in a single input file. Note that the filename cannot contain spaces.

If you wish, the first line of your file can contain a comma-separated list of the available field names; if you choose not to include this line, you must launch your CSV file with a special switch (as described in **Executing the AddImage Command**, page A-8). A CSV file *with* this descriptive line is the default file type expected by the AddImage utility.

The following CSV file would add one image of each type: Class, Topology, Model, and Enterprise. This sample includes the optional initial descriptive line. Note that the line breaks shown here are a publishing limitation only; be sure your file does not break the single line required for each image entry.

CSV Sample File

```
Image type, Image name, Small image file, Large image file,
Image id
```

```
class, Roundabouter, c:\small_icons\round16.bmp,
c:\big_icons\round32.bmp, 711

topology, RoundLAN, c:\small_icons\LAN16.bmp,
c:\big_icons\LAN32.bmp, 712

model, KidMobile, c:\small_icons\car16.bmp,
c:\small_icons\car32.bmp, 713

enterprise, Mom-n-Pop, c:\small_icons\folks16.bmp,
c:\small_icons\folks32.bmp, 714
```

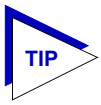

*Although the sample input file shown here assigns a different **Image ID** value to each new image, images of different types can share Image IDs. For example, each image added by this input file could have been assigned the same **Image ID**, as long as that ID value was not being used by another image of the same type.*

Tag/Data Format

To create a file using the multi-line tag/data format, you must place each field name and the value you want to use in its own line. Field names and their values are separated by a colon with a space on either side. The end of each entry is signaled by a separator line containing three dots; any additional blank lines between entries are ignored. Note that the filename cannot contain spaces.

If you choose to use the tag/data format for your input files, you must launch the AddImage utility with a special switch, as described in **Executing the AddImage Command**, page A-8.

The following tag/data formatted file would add one image of each type: Class, Topology, Model, and Enterprise.

Tag/Data Sample File

```
Image type : class
Image name : Roundabouter
Small image file : c:\small_icons\round16.bmp
Large image file : c:\big_icons\round32.bmp
Image id : 711
...
Image type : topology
Image name : RoundLAN
Small image file : c:\small_icons\LAN16.bmp
Large image file : c:\big_icons\LAN32.bmp
Image id : 712
...
Image type : model
Image name : KidMobile
Small image file : c:\small_icons\car16.bmp
Large image file : c:\small_icons\car32.bmp
Image id : 713
```

```
...
Image type : enterprise
Image name : Mom-n-Pop
Small image file : c:\small_icons\folks16.bmp
Large image file : c:\small_icons\folks32.bmp
Image id : 714
...
```

*Although the sample input file shown here assigns a different **Image ID** value to each new image, images of different types can share Image IDs. For example, each image added by this input file could have been assigned the same **Image ID**, as long as that ID value was not being used by another image of the same type.*

Executing the AddImage Command

Once you have created the necessary input files, you can run the AddImage utility via the Run window, as follows:

1. Make sure NetSight Element Manager and all of its components are shut down; the AddImage utility will not run if NetSight Element Manager is running.
2. Select **Start—>Run** from the Windows **Start** button, then enter the following command in the **Open** field:

   ```
   C:\NetSight Element Manager x.x\bin\AddImage.exe [/M]
   [/D] input_filename
   ```

   where:

   /M or **/m** signals that the input file is in multi-line tag/data format

   /D or **/d** signals that the input file is in CSV format, but does not contain the initial descriptive line

   input_filename specifies the path and filename of your input file
3. Click **OK** to launch the utility.

If the tool did not run successfully, an error message will appear, indicating the source of the problem. Note that the tool stops executing at the first error, so if your file contains multiple errors, they will only be detected one at a time. If no message windows appear, the tools have been added successfully.

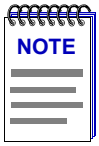

*Once you have added images to the NetSight Element Manager framework, they can only be removed by using the **Restore from new** function in the Data File Manager (described in **Backing Up Data Files** in Chapter 2, **Overview of NetSight Element Manager**). This will restore **all** configurable files to their default values.*

Using the AddDevType Utility

In addition to defining new node class attributes (as described in **Using the AddImage Utility**, page A-5), NetSight Element Manager also allows you to add entirely new node classes. Your new Node Class can use existing Model, Class, Topology, and Enterprise attribute types, or any new ones you have defined via the AddImage utility.NetSight Element Manager provides pre-defined node classes for all of the devices it supports, but if you wish you can use the AddDevType tool to add any additional ones you might need. The AddDevType tool performs the same function as creating a new node class via the Node Class Properties window (as described in **Editing Node Class Properties** in Chapter 4, **List Views**).

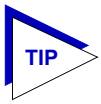

*For more information about creating new image types for use in a new Node Class, see **Using the AddImage Utility** on page A-5.*

As with AddTool and AddImage, to use the AddDevType utility you must prepare an input file in one of two formats — comma-separated variable (CSV), or tag/data format — which contains the following information for each image you wish to add:

Name
Use this parameter to specify the name that will be assigned to your node class.

Model
This field specifies the Model type which applies to the device defined by your new node class. You must select a Model image which has already been defined, and you must be sure to specify the name of the model image using the correct case.

Enterprise
This field specifies the Enterprise type which applies to the device defined by your new node class. You must select an Enterprise image which has already been defined, and you must be sure to specify the name of the enterprise image using the correct case.

sysObjectID
Use this parameter to specify the unique *sysOID* value that has been assigned to the device defined by your new node class. This value must be unique; if you try to define a node class based on a *sysOID* that is already in use, an error message will be displayed and the utility will stop processing the file. If you define a *sysOID* which does not begin with 1.3.6.1.4.1, or one which contains an unrecognized enterprise index, a warning message will appear, but processing will continue.

Classes
This field specifies the Class type or types which apply to the device defined by your new node class. You must select Class images which have already been defined, and you must be sure to specify their names using the correct case. To specify multiple class images, separate the values with commas; for a file in CSV format, enclose the entire attribute value in double quotes.

Topologies

This field specifies the Topology type or types which apply to the device defined by your new node class. You must select Topology images which have already been defined, and you must be sure to specify their names using the correct case. To specify multiple topology images, separate the values with commas; for a file in CSV format, enclose the entire attribute value in double quotes.

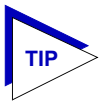

For more information on creating new model, enterprise, class, and topology images, see ***Using the AddImage Utility****, page A-5; for more information on how to view existing images, see* ***Editing Node Class Properties*** *in Chapter 4,* ***List Views****.*

Management Module

Use this parameter to specify the application that will be launched for your new node class when any one of the **Manage** options is selected. If the executable you specify is located with the executables (in the \NetSight Element Manager x.x\bin directory), you need only specify the executable name. If it is located in another directory, be sure to specify the full path.

Management Arguments

You can use this optional parameter to specify any arguments or options that should be passed to the specified command. You should specify any necessary switches in the format required by the application you want to add; if required or allowed, you can also add a specific argument string, or one or more of several available substitutions:

%I	inserts the IP address of the selected node
%C	inserts Community Name associated with the selected node
%O	if a node is selected, inserts its *sysOID*
%N	if a node is selected, inserts its assigned name. Note that, for a device, this will be either the *sysName* value set at the device, or the Name value set via the node's Properties window.
%H	inserts the parent window identifier.
%%	inserts a single percent sign

Working Directory

Use this optional parameter to specify the location of any secondary files that might be needed by the executable, if those files are not located in the same directory as the executable itself. If you choose not to specify a working directory, the directory in which the executable resides will be used by default.

Description

Use this optional parameter to provide a description of the device represented by your new node class. This description will be displayed by default in the Properties window for devices represented by your new node class.

Notes
Use this optional parameter to specify any additional note information you'd like to display by default in the Properties window for devices represented by your new node class.

Creating an AddDevType Input File

You can create your AddDevType input file in one of two formats: multi-line tag/data format, or single-line comma-separated variable (CSV) format. Each method is described below.

Comma-Separated Variable (CSV) Format

To create an input file using the single-line CSV format, you need only string together the fields you wish to use in a single line, with each value separated by a comma. To specify multiple values for a single parameter (for example, multiple Class or Topology images), separate the values with commas, and enclose the entire attribute value in double-quotes. Spaces before and after commas are ignored. If you wish to skip one or more optional fields, be sure to use a comma for a placeholder, or any following fields will be misread. Use one line for each device type you wish to add; extra lines are ignored. Note that the filename cannot contain spaces.

If you wish, the first line of your file can contain a comma-separated list of the available field names; if you choose not to include this line, you must launch your CSV file with a special switch (as described in **Executing the AddDevType Command**, page A-12). A CSV file with this descriptive line is the default file type expected by the AddTool utility.

The following CSV file would create a node class called **Famous_MV**. This sample includes the optional initial descriptive line, and demonstrates using commas as placeholders for the optional fields you do not wish to use and using double-quotes to enclose attributes which contain multiple values. It also uses some of the image types created by the example provided in **Creating an AddImage Input File**, page A-6. Note that the line breaks shown here are a publishing limitation only; be sure your file does not break the single line required for each device type entry.

CSV Sample File

```
Name, Model, Enterprise, sysObjectID, Classes, Topologies,
Module, Arguments, Working directory, Description, Notes

Famous_MV, KidMobile, Mom-n-Pop, 1.3.6.1.4.1.3.3.5,
“Roundabouter, Repeater, Server”, “RoundLAN, HSSI, Ultra
Tech”, c:\dev_manage\manage.exe,,,a very busy LAN bus, Super
crime-fighting device
```

Tag/Data Format

To create a file using the multi-line tag/data format, you must place each field name and the value you want to use in its own line. Field names and their values are separated by a colon with a space on either side. To specify multiple values for a single parameter (for example, multiple Class or Topology images), separate the values with commas. Because each line of data is tagged with its function, you need not use placeholders for optional fields you do not wish to use. The end of each entry is signaled by a separator line containing three dots; any additional blank lines between entries are ignored. Note that the filename cannot contain spaces.

If you choose to use the tag/data format for your input files, you must launch the AddDevType utility with a special switch, as described in **Executing the AddDevType Command**, page A-12.

The following tag/data formatted file would create a node class called **Famous_MV**. This sample demonstrates that optional fields you do not wish to use need not appear in the tool entry; it also demonstrates how to assign multiple values to a single attribute. It also uses some of the image types created by the example provided in **Creating an AddImage Input File**, page A-6.

Tag/Data Sample File

```
Name : Famous_MV
Model : KidMobile
Enterprise : Mom-n-Pop
sysObjectID : 1.3.6.1.4.1.3.3.5
Classes : Roundabouter, Repeater, Server
Topologies : RoundLAN, HSSI, Ultra Tech
Management Module : c:\dev_manage\manage.exe
Description : a very busy LAN bus
Notes : Super crime-fighting device
...
```

Executing the AddDevType Command

Once you have created the necessary input files, you can run the AddDevType utility via the Run window, as follows:

1. Make sure NetSight Element Manager and all of its components are shut down; the AddDevType utility will not run if NetSight Element Manager is running.
2. Select **Start—>Run** from the Windows **Start** button, then enter the following command in the **Open** field:

   ```
   C:\Netsight Element Manager x.x\bin\AddDevType.exe [/M]
   [/D] input_filename
   ```

   where:

   /M or **/m** signals that the input file is in multi-line tag/data format

/D or **/d** signals that the input file is in CSV format, but does not contain the initial descriptive line

input_filename specifies the path and filename of your input file

3. Click **OK** to launch the utility.

If the tool did not run successfully, an error message will appear, indicating the source of the problem. Note that the tool stops executing at the first error, so if your file contains multiple errors, they will only be detected one at a time, each time you run the tool If no message windows appear, the tools have been added successfully.

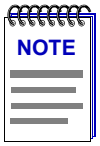

*Once you have added one or more device types to the NetSight Element Manager framework, they can only be removed by using the **Restore from new** function in the Data File Manager (described in **Backing Up Data Files** in Chapter 2, **Overview of NetSight Element Manager**). This will restore **all** configurable files to their default values.*

Viewing the Log File

Each time you run the AddTool, AddImage, or AddDevType utilities, detailed information about the processing of the selected input file and its success or failure will be written to the \NetSight Element Manager x.x\Log\PostInst.log file. You can view this log file via Notepad or any other text-reading application to obtain detailed information about the processing of your input files. Viewing the log file can help you pin down any errors that might be preventing your input files from processing completely; it can also provide a record of the additions you have made to the NetSight Element Manager framework, and the date and time they were made.

Index

Symbols

*.dmf 2-31
*.map files 2-31
*.mmp files 2-31
.dat files 2-31
.MAP files
    inadvertently deleting 6-59
.MRM files 2-31

A

About List Views 4-1
Adding a Device 4-27
Address 5-10, 5-12
Address Mask 2-17
address mask 3-7
Address Range Properties window (Discovery) 3-6
Alarm Threshold 4-34
aligning map elements 6-42
argument substitutions 4-43
Arguments field 4-40
Autofill Community Name 2-15
Autofill IP address 2-14
automatic insertion, when importing maps 6-22
Automatic link updates 6-55

B

Backup Configuration file 2-31
Backup Utility
    performing a backup 2-34
    restoring a backup 2-35
bitmap images 6-9
    as map backgrounds 6-10
    provided 6-11
Building an Argument List 4-42
Buttons 1-2

C

central node database 3-29
Changing Management Application
        Information 4-40
Chassis Manager 2-19, 2-24, 4-35
Choosing a Command (Executable) File 4-41
Class Properties 4-32
Classes 4-14, 4-17, 5-11, 5-14
Classes View 5-5
Clear acknowledged and closed events hourly 2-25
Clear deleted events hourly 2-25
Clear normal and informational events hourly 2-25
Community Name 2-15
Community Name Properties window
        (Discover) 3-11
configuration files, restoring 2-35
Confirm Shutdown 2-15
connections, editing 6-34
Convert window 6-57
copy discovered IP addresses 3-24
    paste into database 3-24
ctron.ini 7-5
Custom Colors 6-10

D

data file backup 2-31
Data File Manager window 2-32
Default Poll Settings (individual node) 2-20
Delete
    Submap icons 6-59
Delete Orphans 4-46
deleting a device 6-58
Deleting a submap icon 6-18
Description 4-30, 4-45, 5-10
device configuration 7-2
device icons, default label 6-26
Device Management 7-1
Device Management options 2-23
Device Properties 4-31
Device Properties, configuring 6-25
devices and servers, moving onto maps 6-23
Directories options 2-21
Discover application 3-1, 3-27
Discover Interval 3-13
Discover Manager 3-1
Discover Manager window 3-14
Discover Manager window, accessing 3-1

Discover Process 3-16
Discover Properties
    Device 3-5
    General 3-4
Discover Properties window, accessing 3-2
Discover Scripts
    deleting 3-18
Discover scripts 3-1, 3-2, 3-19
Discover Scripts, creating and modifying 3-2
Discover Scripts, executing 3-14
Discover Wizard 3-1
    starting 3-19
Discovery Parameters
    IP Address Ranges 3-5
    Subnets 3-7
Display As Icon 6-50
DLM Agent 4-35
DLM Client 4-35
DLM Client (Direct) 4-35

E

embedded object 6-48
embedded objects
    converting 6-56
    updating 6-53
ENETADDR.DEF 7-4
Enterprise 4-14, 5-11
Enterprises 4-16, 5-14
event cache 2-24
    setting the cache size 2-25
Event Log options 2-24
Events 4-14
executable (*.EXE) files 4-40

F

File name 4-41
Find 4-6
Find, in List Views 4-7
Flags, accepted by executables 4-43
folders
    status display 6-6
Font Window 4-14, 5-12
fonts
    configuring for map icons 6-11
    configuring for text elements 6-38

G

General Node Class Properties 4-44
General Properties 4-30
Get MAC address 3-21
Getting Help 1-2
    Using On-line Help 1-2
graphical objects, in maps 6-26
graphical properties, editing 6-35

H

Help Button 1-2
Hint text 2-3

I

ICMP (Internet Control Message Protocol) 7-6
importing maps
    re-discovering devices 6-20
    selecting a position in the hierarchy 6-21
Insert Object window 6-49
Interface List Box 4-32
Interface Number 5-12
Interfaces 4-31, 5-11
Interfaces View 5-5
Interval 4-35
IP Address Autofill 2-17
IP address ranges
    Discovering 3-5
IP addresses
    pasting into database 3-24

L

Label 4-13, 5-10
Large Icon Detail Format 4-6
Large Icon Format 4-5
Launching Device Management 7-1
layering map elements 6-43
Line Properties
    specifying 6-36
linked object 6-48
    updating a link 6-54
linked objects 6-51, 6-52
    changing the source document 6-55
    opening the source document 6-55
    removing the link to the source document 6-56
    updating 6-53
Links window 6-54
locking maps 6-60
Logical Address 4-13
LogSets 7-4

M

MAC address, locating 3-19
MAC search 3-25
manual insertion, when importing maps 6-22
Manual link updates 6-55
Map View 6-2
maps
adding a bitmap background 6-10
adding connections 6-33
adding new 6-7
position in Map View hierarchy 6-8
adding non-manageable graphical objects 6-26
adding shapes and lines 6-34
adding symbols 6-27
and Discover 6-23
and folders, relationships 6-2, 6-6
and Goto symbols, relationships 6-4
and submaps, relationship
restrictions 6-4
and submaps, relationships 6-3
changing 6-17
arranging 6-6
changing the name 6-12
changing the names 6-17
color background 6-10
deleting 6-62
inadvertent deletion, with submap icon 6-18
layering items 6-43
locking 6-60
opening and closing 6-61
placing device icons 6-25
print options 6-11
printing 6-11
saving 2-22, 6-61
setting a grid 6-11
setting display and print options 6-9
status hierarchy 6-3
status reporting 6-1
Method 4-35
Model 4-14, 5-11
Module View window 7-2

N

Name 4-13, 4-30, 4-45, 5-10
Network options 2-15
Node Class page 4-40
node database
polling 2-18
Node options 2-20
Notes 4-30, 4-45, 5-11, 5-12

O

Object 6-48
converting an object's application type 6-56
embedding into a map 6-49
Object Linking and Embedding (OLE) 6-47
object notes 6-39
Object Notes, adding to map elements 6-27
Object Package 6-48
embedding into a map 6-52
Object Packager
updating objects 6-53
Object Type 6-49, 6-55
OLE 6-47
OLE objects, static 6-56
Online Documents 1-2, 1-3
Options window 2-13
Device Management 2-23
Event Log 2-24
Network 2-15
Node 2-20
Polling 2-17
Workspace 2-14

P

Paste Special (No IPs) 6-59
Physical Address 4-13
Ping 3-20
PING (Packet INternet Groper) 7-6
pixels 6-47
Poll Interval 5-12
Poll Method 5-12
Poll Retries 5-12
Poll Settings 4-32
polling
node database 2-19
Polling options 2-17
Print 2-4
Print Preview 2-4, 6-11
Properties 4-32

R

reshaping graphical objects 6-37
resolving host names to IP addresses 3-19
Retries 4-35
Ring Spacing 6-47
rings, moving 6-47

RMFLT.LDB 2-31
rmflt.mdb 2-31

S

Scheduler application
launching from Discover Manager 3-19
scheduling downloads 2-12
scheduling MIB recordings and SETs 2-12
Select All 4-12
Select All Nodes 6-60
Select Orphans 4-46
Services 2-3
Set log size 2-25
Shape Properties
specifying 6-36
Show Splash window 2-14
Show Tooltips 2-14
Small Icon Detail Format 4-6
Small Icon Format 4-5
SNMP 7-2
Source Address naming 7-4
Status 4-12, 5-9, 5-11
status conditions
critical (red) 6-5
normal (green) 6-5
undefined (gray) 6-5
unknown (blue) 6-5
warning (yellow) 6-5
Status Group 4-14, 5-11
Status Groups 4-16, 5-14
status propagation
desktop 4-12, 5-9, 6-5
Submap icons
cutting, pasting, moving, deleting 6-59
submap relationships
creating 6-4
submap relationships, restrictions 6-4
submaps
changing relationships 6-17
creating links 6-13
Subnet Discovery Parameters 3-7
Subnet Properties 3-9, 3-27
Subnet Properties window (Discovery) 3-10
subnet search 3-22
Subnets 4-16, 5-13
System Up Time 4-13, 5-10

T

The All Nodes List View 4-3
The All Nodes Tree View 5-4
The New List View Window 4-10
The New Tree View Window 5-7
Tile Bitmap 6-10
Time of Last Contact 4-13, 5-10
Topologies 4-14, 4-17, 5-11, 5-14
Topologies View 5-5
Topology 5-11
traps 6-5

U

Unnamed PING device 4-29, 6-25
Unselect All 4-12

W

Welcome Wizard 2-15
Workspace options 2-14